

The impact of environment and mergers on the H I content of galaxies in hydrodynamic simulations

Mika Rafieferantsoa^{1,2}, Romeel Davé^{1,2,3}, Daniel Anglés-Alcázar⁴, Neal Katz⁵, Juna A. Kollmeier⁶, Benjamin D. Oppenheimer⁷

¹ *University of the Western Cape, Bellville, Cape Town 7535, South Africa*

² *South African Astronomical Observatories, Observatory, Cape Town 7925, South Africa*

³ *African Institute for Mathematical Sciences, Muizenberg, Cape Town 7945, South Africa*

⁴ *Department of Physics and Astronomy and CIERA, Northwestern University, Evanston, IL 60208, USA*

⁵ *Astronomy Department, University of Massachusetts, Amherst, MA 01003, USA*

⁶ *Observatories of the Carnegie Institute of Washington, Pasadena, CA 91101, USA*

⁷ *University of Colorado, Boulder, CO 80309, USA*

16 September 2015

ABSTRACT

The instantaneous H I content of galaxies is thought to be governed by recent accretion and environment. We examine these effects within a cosmological hydrodynamic simulation that includes a heuristic galactic outflow model that reproduces basic observed trends of H I in galaxies. We show that this model reproduces the observed H I mass function (HIMF) in bins of stellar mass, as well as the H I richness (M_{HI}/M_*) versus local galaxy density. For satellite galaxies in massive ($\gtrsim 10^{12} M_\odot$) halos, the H I richness distribution is bimodal and the median drops towards the largest halo masses. The depletion timescale of H I entering a massive halo is more rapid, in contrast to the specific star formation rate which shows little variation in the attenuation rate versus halo mass. This suggests that, up to the halo mass scales probed here ($\lesssim 10^{14} M_\odot$), star formation is mainly attenuated by starvation, but H I is additionally removed by stripping once a hot gaseous halo is present. In low mass halos, the H I richness of satellites is independent of radius, while in very massive halos they become gas-poor towards the centre, confirming the increasing strength of the stripping with halo mass. Mergers somewhat increase the H I richness and its scatter about the mean relation, tracking the metallicity in a way consistent with it arising from inflow fluctuations, while star formation is significantly boosted relative to H I.

Key words: galaxies: H I, galaxies: formation, galaxies: evolution, methods: N-body simulations, galaxies: ISM, galaxies: mass function

1 INTRODUCTION

The primary components of galaxies are stars and gas. The amount of gas is a vital complement to that of stars as the fuel for new star formation, as well as the repository for the nucleosynthetic by-products of stellar evolution. Therefore, the gas content is a key probe of the life cycle of galaxies. Recently, much progress has been made in understanding the molecular gas content of galaxies throughout cosmic time (e.g. Tacconi et al. 2013) via observations of carbon monoxide (CO) lines and other dense gas tracers. The other major gaseous mass component in galaxies is in the form of atomic neutral hydrogen (H I), which has been more difficult to study to higher redshifts owing to the current sensitivities of radio telescopes. In the nearby Universe, the past decade has seen major H I surveys such

as the H I Parkes All-Sky Survey (HIPASS; Meyer et al. 2004) and the Arecibo Legacy Fast ALFA Survey (ALFALFA; Giovanelli et al. 2005), both able to probe galaxies down to H I masses of $M_{\text{HI}} \approx 10^7 M_\odot$ in a uniform (albeit H I-selected) sample. The GALEX Arecibo SDSS Survey (GASS; Catinella et al. 2010) relaxed the H I selection, and instead selected galaxies based on stellar mass. While they did not probe as deeply in stellar mass (down to $M_* \approx 10^{10} M_\odot$), the lack of H I selection allowed them to assess the biases associated with such a selection technique. With the upgraded Jansky Very Large Array, the CHILES survey aims to probe H I evolution in 21cm emission out to sizeable look back times, to $z \sim 0.45$ (Fernández et al. 2013). Currently, there are two major radio facilities that, as precursors to the

Square Kilometre Array (SKA¹), will probe H I to unprecedented levels both nearby and out to intermediate redshifts. The MeerKAT array in South Africa² and the Australian Square Kilometer Array Pathfinder both have major H I surveys planned, namely the LADUMA³ and DINGO⁴ Surveys, which will probe H I 21cm emission in galaxies out to $z \sim 1$ and $z \sim 0.4$, respectively.

To optimise these investments in observational resources, it is important to place observations of H I within our modern framework for galaxy formation and evolution. Decades ago, Haardt & Madau (2001a) used 324 isolated galaxies and found a more constrained correlation of the optical diameter of a stellar disk with the H I mass than with the morphological type. Even though H I does not directly foster star formation (Kennicutt & Evans 2012), it has been shown that the mass of the stars and the atomic hydrogen (H I) mass are highly correlated (Cortese et al. 2011; Huang et al. 2012). Using HIPASS, Zwaan et al. (2005) measured the H I mass function down to H I masses of $M_{\text{HI}} \approx 10^7 M_{\odot}$, although at low masses the sample is likely biased towards high- M_{HI} systems. They also found dependencies on environment, in the sense that dense regions tend to steepen the H I mass function. An improvement in sensitivity was done by Martin et al. (2010), who observed over 10,000 H I-selected galaxies down to $M_{\text{HI}} \sim 10^6 M_{\odot}$, finding a larger H I cosmic density and a larger value of the characteristic mass, $M_{\text{HI},*}$, from a Schechter function fit.

Interestingly, the galaxy H I and star formation rate observations together showed that the amount of gas in H I is inadequate to yield the observed stellar mass of a given galaxy at a sustained star formation rate (SFR) at any observable redshift (Sancisi et al. 2008). To explain this discrepancy, Hopkins et al. (2008) came up with two possibilities: first, there should be gas replenishment at a rate slightly smaller than the requisite usage (in order for the gas fraction to slowly drop with time), or second, small galaxies are less able to retain their gas. Mass loss from stellar evolution can help the imbalance between the amount of infalling gas and the actual SFR (Leitner & Kravtsov 2011), although not sufficiently in lower mass galaxies and not at epochs much earlier than today. The idea of continued replenishment is also inferred from the evolution of the molecular hydrogen content (e.g. Tacconi et al. 2013). Indeed, current theoretical models of galaxy formation invoke continual gas infall as a driver for star formation (Kereš et al. 2005; Finlator & Davé 2008; Sancisi et al. 2008; Dekel et al. 2009; Davé et al. 2012; Lilly et al. 2013). Hence, it appears that the gas content of galaxies, and particularly the H I content, could provide a probe of gas accretion from the environment around galaxies.

Environment is apparently a key factor in determining the H I content of galaxies. Cortese et al. (2011) found an anti-correlation of H I richness (i.e. H I mass per unit stellar mass) relative to stellar mass, which is valid in moderate to low density environments but highly clustered galaxies tend to be more H I poor (Giovannelli & Haynes 1985;

Solanes et al. 2001). This is often thought to be related to why galaxies in very dense environments also show low specific star formation rates. Pappalardo et al. (2012) found that environment acts to lower the H I content of galaxies. Hughes et al. (2013) also found that low H I objects are mostly found in clustered regions. Data from the ALFALFA $\alpha.40$ catalogue (Haynes et al. 2011) showed that half of the optical sources were observed in denser regions, but only less than one quarter of all the H I detected sources are located in clusters or groups, and most H I-rich galaxies live outside of group environments (Hess & Wilcots 2013).

Mergers, on the other hand, have a significant impact on galaxy evolution. Major mergers, if between gas-rich galaxies, typically induce bursts of star formation (Mihos & Hernquist 1996) that move galaxies well off the so-called star formation main sequence of SFR versus M_* (Noeske et al. 2007), and hence are partly responsible for setting the scatter around the main sequence. Finlator & Davé (2008) argued that departures from the “equilibrium” mass-metallicity relation are driven by stochastic fluctuations in accretion, of which mergers are the most extreme example, and Davé et al. (2013) suggested that this is also reflected in the H I content at a given M_* .

Theoretical work on studying H I is also progressing rapidly. Initially, much of the work utilised semi-analytic models (SAMs) based on prescriptively tying the H I content to the halo mass and merger history. Obreschkow et al. (2009) used the De Lucia & Blaizot (2007) SAM applied to the Millennium simulation, and found that with reasonable parameter choices they could broadly match observations of the H I mass functions for both early and late type galaxies. Lagos et al. (2011) improved on this by adding a prescription to separate H I and H_2 in the GALFORM SAM (Bower et al. 2006) with several different recipes for H_2 , finding substantial differences between such recipes. They also found rapid evolution in H I properties out to $z = 2$. Similar models by Popping et al. (2014) showed that such differences between H_2 recipes are most important in small galaxies. Lagos et al. (2014) further examined the origin of H I in big elliptical galaxies, finding that most of the neutral gas in the elliptical galaxies is produced by radiative cooling from their hot halos. While the large number of free parameters in SAMs makes a unique physical interpretation of the results difficult, nonetheless there is clearly interesting progress being made from SAMs, which are especially useful for making predictions of large upcoming surveys.

Cosmological hydrodynamic simulations have also begun to make predictions for the H I content of galaxies. Popping et al. (2009) presented a simple self-shielding prescription for calculating the H I content of galaxies tuned to match observations of the total cosmic H I content, and showed that such a model produces roughly the correct H I mass function. Duffy et al. (2012) improved on this with a more sophisticated self-shielding model which they applied to the Overwhelmingly Large Simulations (OWLS; Schaye et al. 2010). They obtained good agreement with the observed H I mass function down to $M_{\text{HI}} \sim 10^{9.5} M_{\odot}$, but below this mass they predicted an excess, though this was not very significant since their mass resolution limit was only a factor of several below that; nonetheless, this disagreement mimicked a similar discord in the stellar mass function that is likely a result of their assumed prescription for galactic

¹ <http://www.ska.ac.za>

² <http://www.ska.ac.za/meerkat/>

³ <http://www.ast.uct.ac.za/laduma/Home.html>

⁴ <http://askap.org/dingo>

outflows (Davé et al. 2011b). Davé et al. (2013) presented a model with an improved recipe for galactic outflows that matched the stellar mass function quite well, and applying another improved self-shielding prescription, they were able to also match the H I mass function and H I richness as a function of M_* . This simulation, therefore, provides a plausible model to study how the H I content of galaxies is impacted by other factors in greater detail.

In this paper, we build on the work of Davé et al. (2013) to study the impact of environment and halo mass on the H I content of galaxies and its evolution across cosmic time, particularly focusing on the fluctuations in H I arising from environmental processes in satellites as well as mergers. This paper is structured as follows. In section 2, we begin by reviewing our simulations, particularly the outflow model that is central to matching a variety of observations, as well as our methodology to calculate the H I content of galaxies (2.2) and our methodology for tracking galaxies back in time (2.3). Sections 3 and 4 contain our results. We summarise our results and discuss implications in section 5.

2 METHODS

2.1 Simulations

The main simulation we use here is the same as in Davé et al. (2013), which we briefly review. Using our enhanced version of GADGET-2 (Springel 2005; Oppenheimer & Davé 2008), we run a cosmological hydrodynamic simulation with 512^3 gas particles and 512^3 dark matter particles having masses of $4.5 \times 10^6 M_\odot$ and $2.3 \times 10^7 M_\odot$, respectively, enclosed in a periodic box of $32h^{-1}\text{Mpc}$ comoving on a side, and a comoving gravitational softening length of $1.25h^{-1}\text{kpc}$ (Plummer equivalent). We chose a ΛCDM cosmology consistent with the Wilkinson Microwave Anisotropy Probe results in Hinshaw et al. (2009), namely $\Omega_m = 0.28$, $\Omega_\Lambda = 0.72$, $H_0 = 70 \text{ km s}^{-1} \text{ Mpc}^{-1}$, $h \equiv H_0/(100 \text{ km s}^{-1} \text{ Mpc}^{-1}) = 0.7$, $\sigma_8 = 0.82$, $\Omega_b = 0.046$, and $n_s = 0.96$; these parameters are not far from that favoured by the nine-year WMAP data (Hinshaw et al. 2013) and Planck (Planck Collaboration et al. 2014). GADGET-2 employs entropy-conserving smoothed particle hydrodynamics (SPH), which has some deficiencies related to handling surface instabilities but this makes little difference for the bulk properties of galaxies (Huang et al. 2015, in prep.). We include radiative cooling from primordial (Katz et al. 1996) with additional cooling from metal lines assuming photo-ionisation equilibrium (Wiersma et al. 2009), and track the metallicity in four elements (C, O, Si and Fe) based on enrichment from Type II supernovae, Type Ia supernovae, and stellar mass loss from asymptotic giant branch stars. We will only cursorily be concerned with the metals in this paper, and hence we refer the reader to Oppenheimer (2008) for more details on this aspect. We assume a Chabrier (2003) initial mass function throughout.

The primary distinguishing aspect of our code is the use of a highly constrained heuristic model for galactic outflows. The model we follow is presented in Davé et al. (2013), which utilises outflows scalings expected for momentum-driven winds in sizeable galaxies ($\sigma > 75 \text{ km s}^{-1}$), and energy-driven scalings in dwarf galaxies. The galaxy velocity dispersion of 75 km s^{-1} corresponds to a galaxy baryonic

(star + gas) mass of $M_{\text{gal}} \sim 10^{10} M_\odot$, typically half of which is stars (although gas-poor satellites can be almost entirely stars), at $z = 0$. This mass is lower at higher redshift given that it is a function of $H(z)^{-1}$, as described in Oppenheimer (2008).

In particular, we assume that the mass loading factor (i.e. the mass outflow rate in units of the star formation rate) is $\eta = 150 \text{ km s}^{-1}/\sigma$ for galaxies with velocity dispersion $\sigma > 75 \text{ km s}^{-1}$, and $\eta = 75 \times 150/\sigma^2$ for $\sigma < 75 \text{ km s}^{-1}$. Our previous results have generally favoured these latter scalings for all galaxies, matching everything from intergalactic medium (IGM) enrichment (Oppenheimer & Davé 2006; Oppenheimer et al. 2012) to circumgalactic gas properties (Ford et al. 2013) to galaxy mass-metallicity relations (Finlator & Davé 2008; Davé et al. 2011a) and stellar mass functions (Davé et al. 2011b). However, these prior simulations generally did not resolve dwarf galaxies unlike our current simulation, and hence we found that the steeper scaling of $\eta(\sigma)$ in this regime produced an improved fit to the stellar mass function (Davé et al. 2013). Note that we always assume $v_w \propto \sigma$ for all galaxies, consistent with observations (Martin et al. 2005; Weiner et al. 2009). Interestingly, these scalings are similar to those produced in the fully self-consistent outflow simulations of Hopkins et al. (2014), the Feedback in Realistic Environments (FIRE) suite of zoom simulations (Muratov et al. 2015). The highly simplified physical explanation is that at low masses, supernova energy is sufficient to unbind a sizeable fraction of the gas (Dekel & Silk 1986), but at high masses, additional contributions from momentum input are necessary (Murray et al. 2010; Hopkins et al. 2012).

To implement outflows, we kinetically eject particles from the ISM in order to mimic unfettered escape through ISM chimneys. The mass ejection rate is given by

$$\dot{M}_{\text{wind}} = \eta \times \text{SFR}, \quad (1)$$

where the SFR is computed using a two-phase subgrid ISM model (Springel & Hernquist 2003). If a gas particle has some probability to form stars, it has η times that probability to be ejected in an outflow. We then eject it in a direction given by $\mathbf{v} \times \mathbf{a}$, with \mathbf{v} its instantaneous velocity and \mathbf{a} its instantaneous acceleration. The velocity dispersion of the galaxy is estimated from an on-the-fly friends of friends galaxy finder; see Oppenheimer (2008) for details. After this, we turn off hydrodynamic forces on the particle until it reaches a density that is 10% of the star-forming density threshold, i.e. 0.013 cm^{-3} , or if a timescale corresponding to $1.95 \times 10^{10}/(v_w \text{ km s}^{-1})$ years has elapsed.

This simulation also includes a heuristic model to quench star formation in massive galaxies tuned to reproduce the exponential truncation of the stellar mass function. Star formation in a given galaxy is stopped depending on the quenching probability P_Q given in equation (2), which is a function of the velocity dispersion σ of the galaxy:

$$P_Q = 1 - \frac{1}{2} \text{erfc} \frac{\log \sigma - \log \sigma_{\text{med}}}{\log \sigma_{\text{spread}}} \quad (2)$$

We use $\sigma_{\text{med}} = 110 \text{ km s}^{-1}$ as the velocity dispersion where a galaxy has 50% chance to have its star formation turned off, and $\sigma_{\text{spread}} = 32 \text{ km s}^{-1}$ to describe some scatter between σ and the detailed physics of quenching. Note that this model does not attempt to directly model the physics of quenching,

it is only a way to reproduce the observed mass function at the high mass end, and has virtually no effect below the knee of the mass function. When a galaxy is chosen to stop its star formation, any particle eligible for star formation first has its quenching probability assessed, and if it is selected for quenching then it is heated to 50 times the galaxy’s virial temperature, which unbinds it from the galaxy. As discussed in Davé et al. (2013), and as we will demonstrate later, this quenching model does not substantially impact the overall H I content of galaxies, though there are some important environmental effects. It is also important to note that, owing to being tied to a threshold in σ , this quenching prescription really only affects massive *central* galaxies; the quenching of *satellite* galaxies occurs owing primarily to tidal and hydrodynamic stripping effects that are self-consistently modeled in the simulation.

To examine dependencies on our feedback assumptions, we will use another model without quenching, namely the “vzw” or momentum-driven wind galactic outflow model. This model, fully described in Davé et al. (2011a), has two main differences relative to the primary model used in this work. First, the vzw model does not include the quenching prescription as in our ezv model. Second, the mass loading factor η is inversely proportional to the velocity dispersion σ regardless of the size of the system. Hence, we would only expect a difference at the low σ end ($\sigma < 75 \text{ km s}^{-1}$) owing to the ezv model adopting a steeper outflow scaling ($\eta \propto \sigma^{-2}$) at these smaller σ . For massive galaxies, we will use the comparison between the two models to distinguish the effect of the star formation quenching prescription, which only operates at high σ where the two models have the same mass loading factor η dependency.

Our analysis will only consider galaxies that contain at least 64 star particles having a respective stellar mass of $1.45 \times 10^8 M_\odot$. With this limit, the simulation produces 3,732 galaxies, among which 2,607 and 1,125 are central and satellite galaxies, respectively. At the low masses, it was shown in Davé et al. (2013) that galaxies are most H I rich, with a gas content that is typically at least as much as their stellar mass (in accord with observations: e.g. Walter et al. 2008; Rupen & LITTLE THINGS Team 2010). Given that a gas particle can spawn to 2 star particles⁵, our stellar mass limit corresponds to 32 gas particle mass. This is probably suboptimal for fully resolving ram pressure stripping processes, as the idealised tests of ram pressure stripping in hot cluster gas by McCarthy et al. (2008) suggest that significant non-convergence appears when the number of gas particles drops below several hundred. Hence our results for galaxies with $M_* \lesssim 10^9 M_\odot$ may be considered preliminary subject to more detailed mass resolution convergence studies. Nonetheless, the results for our lowest mass galaxies tend to mostly lie on an extrapolation from higher-mass galaxies, and hence the qualitative trends we identify in this work are likely to be robust. Spatially, even our smallest galaxies have a typical H I-weighted radius of $\sim 10 \text{ kpc}$, well above our spatial resolution. The tests for resolution convergence were mainly done by looking at the differences between our high

resolution simulation and another lower resolution simulation with only 2×256^3 particles, i.e. eight times less particles. We find that the two mass functions are very similar above the mass resolution limit, which implies good convergence. The H I richness versus stellar mass, however, has noticeable differences between the two simulations, with the high resolution simulation being more H I poor by $0.1 - 0.2 \text{ dex}$. This discrepancy owes to the contributions of both an increased stellar mass and a decreased H I mass for a given halo mass at higher resolution, since higher resolution simulations are more efficient in cooling the gas and hence in forming stars. Despite these relatively minor differences, the resolution convergence is sufficient not to affect our main conclusions. For more detailed tests of resolution convergence, we refer the reader to Davé et al. (2013).

Each halo is identified using a spherical overdensity (SO) algorithm (Kereš et al. 2005), which takes each galaxy centre and expands a sphere around it to enclose an average overdensity given by

$$\delta = 6\pi^2(1 + 0.4093(1/f_\Omega - 1)^{0.9052}) - 1 \quad (3)$$

with

$$f_\Omega = \frac{\Omega_m(1+z)^3}{\Omega_m(1+z)^3 + (1 - \Omega_m - \Omega_\Lambda)(1+z)^2 + \Omega_\Lambda}. \quad (4)$$

This corresponds to an overdensity of $\delta \simeq 110$ at $z = 0$. In our simulation volume, the most massive halo at redshift $z = 0$ is $M_{\text{halo}} = 1.32 \times 10^{14} M_\odot$; hence our environmental study focuses on the field to moderate-sized group environment.

2.2 Computing the H I content

We identify galaxies using Spline Kernel Interpolative Denmax⁶ (SKID) as bound collections of stars and star-forming gas. We associate each galaxy with a halo identified using a spherical overdensity algorithm, as described above. The central galaxy is taken to be the most massive galaxy in the halo, and we consider all the others as satellites. While SKID adequately captures the stellar and molecular mass, significant amounts of H I can be present outside of the star-forming gas. We compute the H I content of the simulated galaxies following the methodology described in full detail in Davé et al. (2013); we review the main points here as well as modifications to the previous modeling.

First, we consider all particles within a sphere of a radius given by the outermost particle in the SKID galaxy, and associate all the H I to that galaxy. For close pairs of galaxies, we assign the particles to the galaxy with higher gravitational force to avoid double-counting H I. For each gas particle, we calculate the neutral fraction based on the assumption that the given particle is a sphere with a kernel density profile given by the SPH kernel, and that it is bathed in radiation from the metagalactic radiation field given by Haardt & Madau (2001b). We then integrate the column density inwards from the particle surface until it reaches a threshold column density where the particle becomes sufficiently neutral. The mass fraction within this radius is then considered to be 90% neutral (since some mass in the outskirts of galaxies remains ionised), while the

⁵ This number is chosen as a compromise between desiring a gradual buildup of stellar mass, versus not introducing too many new particles that slows the computation.

⁶ <http://www-hpcc.astro.washington.edu/tools/skid.html>

mass outside this radius has the optically-thin ionisation fraction. We refer to this as the “auto-shielding” approximation. and Davé et al. (2013) showed that it yields good agreement with the full radiative transfer simulations of Faucher-Giguère et al. (2010). We tested the same method by using a sphere of radius $1.5\times$ higher than that given by the outermost particle in the SKID galaxy and ended up with a $\lesssim 5\%$ typical increase in the total H I mass.

Compared to Davé et al. (2013), we have changed the column density threshold within which we consider the gas to be fully neutral in the auto-shielding calculation. Davé et al. (2013) utilised $N_{\text{HI}} = 10^{17.2} \text{ cm}^{-2}$, which corresponds to an optical depth of unity to ionising radiation. However, given that the ionisation level of an optically-thin gas bathed in metagalactic flux is much smaller than unity, typically $10^{-5} - 10^{-6}$ (e.g. Davé et al. 2010), an optical depth of unity will still not produce a substantial H I fraction. Instead, we need an optical depth that will result in an H I fraction of 0.5. If we assume that the ambient gas has an ionisation fraction of 10^{-5} , we require that the optical depth be 10.8 instead of unity. The resulting column density threshold thus rises from $1.6 \times 10^{17} \text{ cm}^{-2}$ to $1.7 \times 10^{18} \text{ cm}^{-2}$. We, therefore, use this latter column density threshold to calculate the radius within which the gas is assumed to be fully neutral. We have checked that the agreement with the radiative transfer simulations of Faucher-Giguère et al. (2010) remains acceptable, and in fact is slightly improved. In §3.1 we will compare our results to that using the old threshold value. If such a radius does not exist (typically for particles with density $n \lesssim 10^{-2} \text{ cm}^{-3}$), then the particle is fully optically thin and its ionisation fraction is calculated assuming ionisation balance with the metagalactic flux.

Next, we must also determine how much of this shielded gas is molecular. For this, we employ the H_2 formation model of Krumholz & Gnedin (2011), which we compute for each particle. They consider radiative transfer coupled with the formation and dissociation balance of the molecular hydrogen in a steady state molecular cloud. The resulting H_2 fraction is given by

$$f_{H_2} \simeq 1 - \left(\frac{3}{4}\right) \frac{s}{1 + 0.25 \times s} \quad (5)$$

where s is a term dependent on the formation rate of the molecular hydrogen from dust grains (which depends on metallicity), the dust cross section per H nucleus, and the ambient intensity of the ultra-violet radiation field. Note that in Davé et al. (2013) the default model was the observationally-constrained pressure law of Leroy et al. (2008). We use Krumholz & Gnedin (2011) here because we want to study the evolution of H I to higher redshifts, where the metallicities can significantly deviate from the solar abundances typical in the Leroy et al. (2008) data; the Krumholz & Gnedin (2011) model accounts for such metallicity dependencies in H_2 formation. Davé et al. (2013) demonstrated that at $z = 0$ this makes essentially no difference to most H I properties, but at higher redshifts there are more substantial differences. Briefly, high redshift galaxies have lower metallicities, resulting in less molecular gas and hence an increased H I mass fraction. Davé et al. (2013) showed that at $z = 3$ the HIMF using the Krumholz & Gnedin (2011) model results in a mass function that is about two times higher than that derived using the

Leroy et al. (2008) prescription. Nonetheless, in this work we mostly focus on lower redshifts where the differences are minimal.

With this prescription, we separate each gas particle into a neutral, molecular, and ionised component. The H I mass of the galaxy is then the sum of the neutral mass in all its associated particles, and likewise for the molecular mass.

2.3 Tracking progenitors

In this paper, we will be interested in tracking the evolution of H I in individual galaxies back in cosmic time. We do this by associating a given galaxy at $z = 0$ to its most massive progenitor at all of our previous 134 outputs back to $z = 30$. We define the most massive progenitor as the galaxy (identified by SKID) at an earlier output containing the largest number of star particles in common with the galaxy at $z = 0$. We choose the limiting number of particles for a group to be considered as the galaxy progenitor to be 32 particles, which is below the resolution limit. This choice does not affect the resolved progenitor history, but allows the history to be extended somewhat further back in time.

From this, we can identify major mergers by a simple prescription, following Gabor & Davé (2012). For an assumed major merger ratio r , we search the stellar mass growth history for jumps in excess of $1/(1+r)$. The minimum value of r for which continual infall and a merger could be distinguished depends on the redshift and the time between successive outputs, but it is generally quite far below the minimum ratio of 1 : 3 that we use in the analysis, i.e. our outputs are frequent enough, every 100 to 300 Myr, such that no galaxy would ever grow by 33% just from *in situ* star formation between outputs alone. There are, however, occasional complications with this approach. First, a star particle might, at a given output, be attributed to no group. Second, two groups brought close along their path (without merging) could be regarded as one group by SKID, but would later “unmerge”. To address the first problem, we look back in time until the star belongs to a galaxy, and then attribute that star to the descendent of that galaxy. To address the second problem, we check which galaxies at the previous output are composed of two separate galaxies. We track those merging units until every star in both galaxies is assigned to a single galaxy. From then on, we assign the stars located in the smaller group to that single galaxy. We have found that this robustly identifies merger events despite the dynamical nature of the merger encounter.

3 H I AND ENVIRONMENT AT $z=0$

3.1 H I mass function

The most basic statistical property of H I galaxies is the H I mass function (HIMF). Davé et al. (2013) showed that the simulation we use here, together with their prescription to compute the H I content, yields an H I mass function that is in good agreement with observations down to the lowest resolvable masses ($\sim 10^8 M_\odot$). This is a non-trivial success that has been difficult to achieve in simulations. Recent SAMs have done better but still show an excess at $M_{\text{HI}} \sim 10^9 M_\odot$ (Lagos et al. 2014). In this section

we test our HIMF in more detail, in particular separating it into bins of stellar mass to compare with recent data from Lemonias et al. (2013).

Figure 1 shows the HIMF of our galaxies (black line with blue circles), down to an H I mass similar to the stellar mass resolution limit, namely $M_{\text{HI}} = 1.45 \times 10^8 M_{\odot}$. Note that our simulation does not have a formal H I mass resolution limit, only a stellar mass resolution limit. Here we leverage the fact that, near this resolution limit, the H I and stellar masses are typically comparable, and hence we quote our HIMF down to this mass limit in H I as well. However, this means that we may miss some galaxies with large H I fractions below our stellar resolution limit. We found that reducing the stellar mass limit by a factor of two results in only a $\sim 15\%$ increase in the H I mass function at $M_{\text{HI}} \approx 1.4 \times 10^8 M_{\odot}$. This suggests that our predicted HIMF faint end is not very sensitive to our stellar mass cut excluding low- M_* , large M_{HI} galaxies.

Since our N_{HI} self-shielding threshold is altered from that assumed in Davé et al. (2013), we compare our HIMF to theirs shown as the dotted green line. Overall, this change is consistent with a $\sim 30\%$ shift towards lower M_{HI} (i.e. leftwards) versus the Davé et al. (2013) HIMF, since it results in a lower fraction of each gas particle being neutral. Consequently, this produces a larger change in the amplitude at the high-mass end. Our new HIMF is in better agreement with the observed HIMF from the ALFALFA survey (dashed line; Haynes et al. 2011), modulo the concerns discussed above regarding missing low- M_* galaxies with very high M_{HI} .

We further divide our HIMF into bins of 0.5 dex in stellar mass M_* , shown as the coloured lines, from $M_* = 10^8 M_{\odot} - 10^{12} M_{\odot}$. This shows how the HIMF is comprised of galaxies with different stellar masses. Since M_{HI} correlates with M_* , there is a general trend that the HIMF shifts towards higher M_{HI} at higher M_* . There is also a slight trend for an increased width of the HIMF as one moves towards higher M_* , reflective of the fact that low- M_* galaxies have more uniformly high H I richness, while more massive galaxies can have a wider range of H I content. To quantify this wider range in M_{HI} at higher M_* , we perform Gaussian fits to the HIMF for each stellar mass bin. We compute the mass range $\Delta \log M_{\text{HI}}$ centred at the mean value of the Gaussian fit that encloses 90% of the galaxies in each M_* bin, which is indicated as the “w90” value. We then plot “w90” versus M_* in Figure 2, with the same colour scheme as in Figure 1. w90 steadily increases from 1.03 dex for the lowest M_* bin to 1.76 dex for the highest M_* bin, roughly following the relation $\Delta \log(M_{\text{HI}}/M_{\odot}) \simeq 0.23 \log(M_*/M_{\odot}) - 0.9$.

If there was perfect correlation between M_{HI} and M_* , then the HIMFs separated into M_* bins would have no overlap. Hence, the fact that the HIMF spreads over a larger range than the bin size in M_{HI} is a measure of the scatter between M_* and M_{HI} . This scatter is a key indicator of how the H I content of galaxies varies with environmental influences. For instance, Moran et al. (2012) showed that galaxies with enhanced total M_{HI} tend to have the excess H I in their outskirts, accompanied by a metallicity drop, indicative of recent accretion, although Wang et al. (2014) find that roughly half of H I rich galaxies have the excess predominantly in their centre, and not in the outskirts. One possible physical interpretation is that the spread between




Figure 2. The mass range $\Delta \log M_{\text{HI}}$ enclosing 90% of the galaxies in the HIMF (w90) versus the mid-value of the M_* bin. Each stellar mass-binned HIMF shown in Figure 1 is fit with a Gaussian and then w90 is computed. The colour scheme is the same as that in Figure 1. w90 steadily increases from 1.03 dex at the lowest M_* bin to 1.76 dex at the highest M_* bin. The black line shows a log-linear fit with a slope 0.23.

M_{HI} and M_* tracks the frequency and amount of recent accretion (or lack thereof) that gives rise to variations in H I without accompanying immediate variations in M_* .

To test this aspect of our simulations, we can compare our results to the data from the GASS survey⁷ separated into stellar mass bins by Lemonias et al. (2013). The GASS survey is stellar mass-selected down to $M_* = 10^{10} M_{\odot}$, and hence provides a fair comparison to our stellar mass-selected simulated galaxy sample, at least down to their completeness limit.

The right panels of Figure 1 show the comparison to the data from Lemonias et al. (2013), shown as the green diamonds for the H I-detected galaxies only, and as black circles for all galaxies including those without detected H I placed at their upper limit in H I, in three bins with widths of 0.25 dex in M_* . For a fairer comparison, we similarly adjusted the H I mass of the simulated galaxies according to the GASS data limits used in Lemonias et al. (2013), namely we set a minimum mass of $M_{\text{HI}}/M_* = 1.5\%$ for galaxies with $\log_{10}(M_*/M_{\odot}) > 10.5$ and $\log_{10}(M_{\text{HI}}/M_{\odot}) = 8.7$ for galaxies with $\log_{10}(M_*/M_{\odot}) < 10.5$. Hence the simulation points should most directly be compared to the open black circles; the green diamonds are shown so the reader can assess the contribution from non-detected galaxies to the observed sample. Error bars on the HIMF indicate the cosmic variance in each M_{HI} bin; they are computed as the scatter of the HIMFs amongst the octants of the cubical simulation volume.

At low H I masses, the observation data (black circles) fall within the simulation predictions’ (blue line) cosmic variance. At high masses, however, our simulated galaxy population does not extend up to the largest H I galaxies, showing that our model does not produce galaxies at $M_* \gtrsim 10^{10} M_{\odot}$ with very high H I content. This may partly owe to our quenching prescription that, in some cases, starts to affect galaxies at these stellar masses; we will demonstrate later that once galaxies enter the mass regime where our quench-

⁷ <http://www.mpa-garching.mpg.de/GASS/>

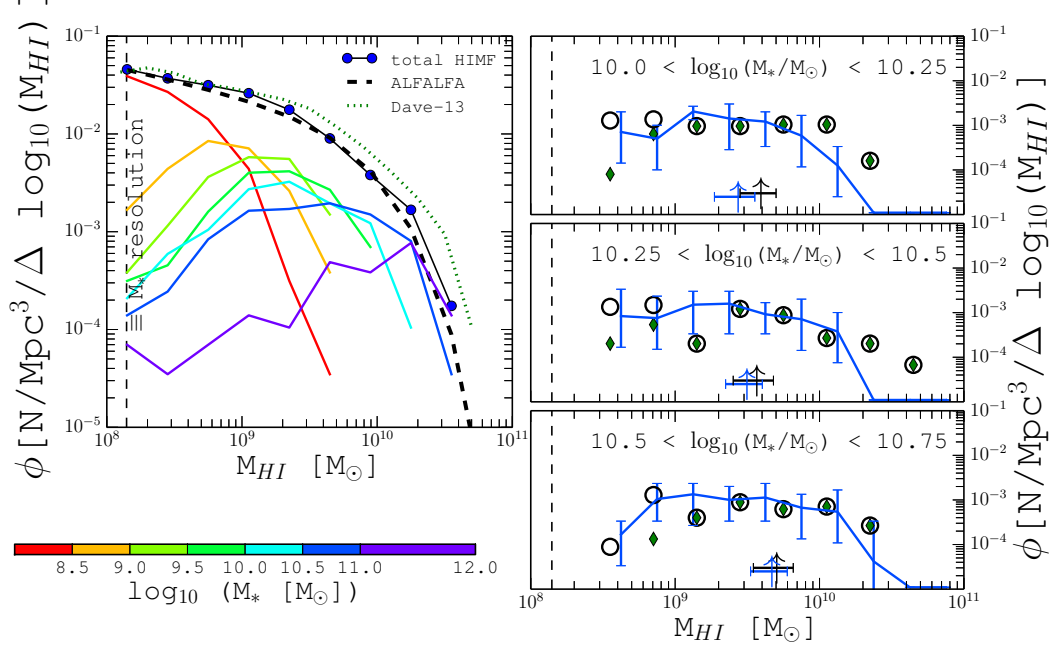


Figure 1. The H I mass function at $z = 0$ for the population of galaxies generated by the simulation described in section 2. *Left panel:* We plot our total HIMF as the black line (with blue circles), the simulated HIMF of Davé et al. (2013) as the dotted green line, and the HIMF of galaxies from the ALFALFA H I survey (Haynes et al. 2011) as the black dashed line. Contributions to our simulated HIMF from bins of 0.5 dex in stellar mass are indicated by the colored lines shown by the color-bar for the mass range $M_* = 10^8 - 10^{12} M_\odot$ (the highest mass bin covers 1 dex owing to the small number of galaxies in this mass range). The vertical dashed line shows the stellar mass resolution limit for reference; there is no formal H I mass resolution limit, but Figure 3 shows that the H I and stellar masses are comparable in galaxies of this mass. *Right panel:* Comparison between our simulated galaxy HIMF (blue lines; where error bars indicate the cosmic variance in each M_{HI} bin) and the observational data of Lemonias et al. (2013) (green diamonds: H I detected sample; open black circles: full sample with non-H I-detected galaxies included at their upper limits) for three different bins in stellar mass. The upward arrows on the bottom of each right panel indicate the weighted average of the H I mass with errors computed as described in the text (black: observed sample; blue: simulation prediction).

ing starts to become important, our simulation does not reproduce the observations quite as well, suggesting that our quenching model is not fully compatible with the H I data. On the other hand, the agreement is if anything somewhat better at the highest M_* values, suggesting perhaps that this is a more fundamental issue with our simulation.

Near the bottom of the right panels of Figure 1, we show the weighted average of the H I mass in each stellar mass bin as the upward arrows. We use the normalised number density of galaxies ($\phi_i / \sum \phi_i$) to determine the central value of each H I mass bin. The errors are also weighted by the normalised number density of galaxies, taking the error for each bin as the bin width in H I mass. The colours correspond to the various data points; we only show the values for the data and models with H I-undetected galaxies set at the detection limit (black and blue upward arrows, respectively). The data are more consistent with our simulations in the higher stellar mass bins. The differences at low masses mainly owe to the fewer high H I mass galaxies in the lower stellar mass bin, and for this reason, the weighted H I mean is lower for the simulated galaxies. The difference between the means decreases from $\sim 40\%$ to $\sim 10\%$ in the highest stellar mass bin.

This comparison suggests that our simulations are generally reproducing the scatter in the relationship between

M_{HI} and M_* , at least down to $M_* \approx 10^{10} M_\odot$. To quantify this further, we plot in Figure 3 H I richness (M_{HI}/M_*) as a function of M_* . We subdivide it into central (red) and satellite (blue) galaxies. The H I-poor galaxies dominate the satellite population especially at low stellar masses, while the central galaxies drive the shape of the relationship between H I-richness and M_* because they are more numerous. In other words, a similarity can be seen between the black line and the red dashed line, where the maximum difference is $\sim 60\%$ at $M_* = 1.45 \times 10^8 M_\odot$. We will discuss the satellite population further in §3.2. The grey hatched area represents the 1σ uncertainty of the simulated sample falling in the same range of stellar mass as the GASS data. This is identical to Figure 4 in Davé et al. (2013), except that we are now employing our new N_{HI} self-shielding limit as described in §2.2.

We show the median of the GASS data (open circles) from their third data release (Catinella et al. 2013) with the error bars representing the 1σ uncertainty in the median computed using jackknife resampling over 8 subsamples. The simulations and the observations are consistent at the $\lesssim 2\sigma$ level at lower masses, but at higher stellar masses ($M_* \gtrsim 10^{10.5}$) we see that the simulation (black line) deviates from the observations (open circles). This deviation owes to our quenching model, which stops the stellar mass




Figure 3. The H I richness (M_{HI}/M_*) vs. M_* of $z = 0$ galaxies from simulations and observations. Red points are central galaxies, blue points are satellites. The galaxies grouped near the bottom of the plots contain galaxies with $M_{\text{HI}}/M_* \in [0, 10^{-3}]$, with their distribution peaking at $\sim 10^{-6}$. The black line is the running median of the simulated galaxy H I richness, with the 1σ uncertainty indicated by the grey hatched area on the massive end (overlapping with GASS). The data from GASS DR3 (Catinella et al. 2013) are shown as the empty circles with error bars representing 1σ uncertainties of the medians. The non-detected sources in GASS are included at their upper limit. ALFALFA data (Huang et al. 2012) are presented as black squares. The striking difference between ALFALFA and GASS owes to the selection criteria, as described in Catinella et al. (2010). The stellar mass selection of GASS is more directly comparable to our simulated sample, which is consistent with GASS observations within the uncertainties. The lower panel is the log amplitude of the scatter around the central galaxies; the scatter generally increases with mass.

growth, but the galaxies can still accrete some gas. Qualitatively, this suggests that our quenching model needs to be more effective at preventing cold gas accretion. We also show the results from the ALFALFA survey (black squares; Huang et al. 2012) for reference, but these are not for direct comparison to our models. These values are higher than those seen in the GASS survey since ALFALFA is H I-selected and hence picks out higher- M_{HI} galaxies at any given mass. To compare to ALFALFA we would have to mimic the volume and selection function of that survey, which we leave for future work⁸. Hence, while the ALFALFA HIMF is appropriate to compare to our models (under the reasonable assumption that ALFALFA has appropriately accounted for its selection volume), the H I richness from that

survey is biased high (Catinella et al. 2010). Thus our predicted H I richness should not be directly compared to the ALFALFA data. Nonetheless, it is interesting that our models follow the shape of the ALFALFA H I richness vs M_* to low masses, offset by a fixed factor that is similar to the offset between GASS and ALFALFA in their overlapping M_* range.

In summary, our simulated HIMF agrees very well with ALFALFA observations, particularly with our more physically-motivated criterion for the column density where auto-self-shielding becomes important. Binning the HIMF into M_* bins, for non-quenched galaxies, our simulation agrees well with observations from the GASS survey, suggesting that the scatter in M_{HI} versus M_* also agrees with observations, although the simulations do not produce the most H I-rich galaxies at lower stellar masses. For quenched galaxies (or galaxies on the verge of quenching), we likely require a more sophisticated quenching prescription to match the data at the low- M_{HI} end, and this could potentially provide a new and interesting constraint to test our ongoing improvements to our quenching model (e.g. Gabor & Davé 2015). This is also evident in our H I richness predictions, which match well with stellar mass-selected observations up to the masses where once again our quenching model might be inadequate. Modulo these small discrepancies, the general agreement in the scatter between M_{HI} and M_* is encouraging and the first time it has been demonstrated for cosmologically-based H I models of any type.

3.2 Satellite galaxies

Once a galaxy falls into another galaxy’s halo, a number of physical processes can act to remove or deplete its gas. These include tidal stripping, ram pressure stripping (e.g. Gunn & Gott 1972; van Gorkom et al. 2003), viscous stripping (e.g. Marcolini et al. 2003), induced star formation from encounters with other galaxies (harassment; Moore et al. 1996), and starvation of inflow to sustain star formation. Hence the behaviour of satellites within a halo is expected to differ from that of the central galaxy. These various physical processes are expected to leave different signatures on the H I content, based on how much ambient gas is within the central galaxy’s halo. Ram pressure stripping scales as $\propto \rho v^2$ (Gunn & Gott 1972), where ρ is the ambient gas density and v is the relative velocity of the satellite, and hence is effective in high velocity dispersion halos with significant halo gas. Induced star formation relies on the presence of other satellites to stimulate a starburst (from tidal interactions; Li et al. 2008). Starvation in satellites occurs when its own hot gaseous halo gets stripped, and its feeding streams have been disconnected since they now head towards the central galaxy. The H I gas, generally being more loosely bound and hence more easily susceptible to stripping and heating than the molecular gas and stars, provides a unique and interesting test bed to study the importance of these environmental processes over a wide range of halo masses. In this section we focus on the properties of centrals and satellites as a function of halo mass.

Figure 4 shows the H I mass (upper panel) and H I richness (lower panel) as a function of halo mass for galaxies in our simulations. Central galaxies are shown as red points, and satellites are indicated by blue crosses. The data points

⁸ We note that this has caused some confusion in the literature. For instance, the Illustris simulation (Vogelsberger et al. 2014) agrees well with the ALFALFA H I richness data, but this is using all their galaxies without mimicking the ALFALFA selection.

along the bottom of each panel correspond to galaxies with H I mass (or H I richness) below the y-axis limit, and they are given some artificial scatter for visibility. The magenta lines and blue dashed lines show the running median values for centrals and satellites, respectively. The green lines show the median values for the central galaxies within the stellar mass range of the satellite galaxies in the same halo mass bin; this specifically isolates the difference between centrals and satellites owing to halo mass, in a stellar mass-matched sample. Histograms in the upper panel (red: centrals; blue dashed: satellites) show the distribution of all resolved galaxies ($M_* \geq 1.45 \times 10^8 M_\odot$) in terms of halo mass. Histograms in the bottom panel indicate the distribution of galaxies with $M_{\text{halo}} > 10^{12} M_\odot$ (along with our resolution limit of $M_* \geq 1.45 \times 10^8 M_\odot$). The amplitudes of the histograms in each panel are arbitrary, as these are purely for comparing the abundance of centrals versus satellites.

The top panel shows that the H I mass of centrals increases steadily with halo mass, though there is a marked transition at $M_{\text{halo}} \gtrsim 10^{12} M_\odot$ above which the relationship is less steep. This trend mimics the drop in sSFR around this halo mass (e.g. Davé et al. 2011b; Salim et al. 2007), which arises owing to a combination of the increased presence of a hot gaseous halo along with quenching feedback. Meanwhile, the H I mass in satellites is roughly independent of halo mass, rising with centrals at low masses but decreasing in more massive halos. The histogram along the bottom shows that satellites strongly increase in relative abundance at larger halo masses, as expected from halo occupation distribution statistics (e.g. Berlind et al. 2003). In addition, the upper panel also shows the median value of the H I masses of the central galaxies having the same range of stellar mass as the satellites in the same halo mass bin (green line). We can only compute this quantity for $M_{\text{halo}} \lesssim 10^{13} M_\odot$, because we chose to select central galaxies within $\sim 2\sigma$ of the satellite stellar mass range, and the central galaxy stellar masses lay outside of this range in the highest halo mass bins. Despite this limitation, we can still see that the effect of halo mass on the gas content of satellites is still visible and seems to become more prominent beyond $M_{\text{halo}} \simeq 10^{12} M_\odot$.

For H I richness (Figure 4, bottom panel), the centrals and the satellites follow a similar trend in the low mass regime, with the satellites having slightly higher H I richness owing to their typically lower M_* . Hence at $M_{\text{halo}} \lesssim 10^{12} M_\odot$, it seems that satellites' H I content is essentially unaffected by being within the halo of another more massive galaxy.

The situation is different for $M_{\text{halo}} \gtrsim 10^{12} M_\odot$. Here, the median satellites H I richness drops dramatically, and even though their masses are lower (and hence by the overall correlation should have higher H I richness), their H I richness is significantly below that of centrals. Looking more closely, this drop is driven by a rapidly growing population of satellites with essentially no H I, whereas the satellites that have H I tend to lie slightly above the centrals following the trend at lower halo masses. The blue histogram (dashed) along the right axis shows this bimodality for galaxies with $M_{\text{halo}} > 10^{12} M_\odot$. This strong bimodality, analogous to that in a galaxy colour-magnitude diagram, indicates that the H I in satellites is fairly rapidly removed once it enters a massive halo. While the lower peak is enhanced in amplitude since we have placed all H I-poor galaxies at $\sim 2 \times 10^{-4}$, nonetheless




Figure 4. The H I mass (top) and H I richness (bottom) of simulated galaxies as a function of halo mass at $z = 0$. Centrals are shown in red (points), satellites in blue (crosses), and the magenta lines and blue dashed lines correspond to their running median values respectively. The green lines show the median values of the central galaxies having stellar masses within the stellar mass range of the satellite galaxies in the same halo mass bin (“ M_* -matched” centrals). Histograms in the upper panel show the distribution of halo masses for all central (red line) and satellite (blue dashed line) galaxies. Histograms in the lower panel show the distribution of H I richness for centrals (red line) and satellites (blue dashed line) with $M_{\text{halo}} > 10^{12} M_\odot$. Environmental effects become very prominent for satellite galaxies with $M_{\text{halo}} \gtrsim 10^{12} M_\odot$.

there is clearly an emerging population of low-H I galaxies that is distinct from the extrapolated trend from lower mass halos, regardless of the binning.

Note that we only considered $M_* \geq 1.45 \times 10^8 M_\odot$ galaxies, which could in principle underestimate the median H I richness of satellites if this mass resolution limit excludes a population of lower stellar mass satellites with very high M_{HI}/M_* . However, for this to happen, there would suddenly have to be many H I-rich galaxies in massive halos just below our M_* resolution limit. This seems rather unlikely, given that in Figure 5 we will show that virtually all low-mass satellites in massive halos are H I-poor. Also note that there are some H I-poor central galaxies with low halo masses. As discussed and quantified in Gabor & Davé (2015), these are likely satellite galaxies whose orbits take them just outside their massive host halo, and thus are identified as centrals by our spherical overdensity halo finder. The contribution of




Figure 5. The H I richness of satellite galaxies versus their stellar mass at $z = 0$. The galaxies are colour-coded by their respective halo mass as shown in the colour bar. The lines represent the median values for each halo mass bin. The points at the bottom of the figure show galaxies with H I richness below 10^{-3} . This plot shows that the host halo mass starts to influence the majority of satellites at $M_{\text{halo}} \gtrsim 10^{12} M_{\odot}$.

these galaxies is minor, but had we identified these as satellites of massive halos, this would strengthen the bimodality. We will examine the role of satellite orbits in stripping and bimodality in a future paper.

We show in Figure 5 the relation between H I richness and stellar mass for satellite galaxies within different halo mass bins differentiated by their colours: from black for $\log(M_{\text{halo}}/M_{\odot}) \sim 11$ to yellow for $\log(M_{\text{halo}}/M_{\odot}) \sim 14$. From the figure, we can see that starting from $M_{\text{halo}} \sim 10^{12} M_{\odot}$ (as mentioned before), there is a clear evidence for a growing population of gas-poor satellite galaxies. The H I-poor satellite population (shown by the data points near the bottom of the plot) is dominated by galaxies located in more massive halos. Small satellites (low stellar mass) get their gas content easily stripped owing to less gravitational pull binding the gas to the galaxy, leading to the higher number of galaxies at lower stellar masses in the highest halo mass bins. This is an interesting prediction, but observation are not yet capable of detecting the H I content of such small galaxies in dense environments expected to be dominated by hot halo gas. For instance, the upper limits in Cortese et al. (2011) are well above what we predict for H I content. Meanwhile, the LITTLE THINGS⁹ survey observes only very nearby dwarf galaxies, which do not lie in sufficiently massive halos.

The oft-noted emergence of a hot gaseous halo around that mass scale (e.g. Kereš et al. 2005; Gabor & Davé 2012) hints at a connection between the presence of a hot halo and the removal of H I, which favours a removal mechanism associated with gas stripping. However, there are also many more satellite galaxies in such massive halos that could tidally strip or harass satellites, so the connection is not completely certain. Nonetheless, as shown in Cunnama et al. (2014) from hydrodynamic simulations, H I-poor satellites

generally feel a ram pressure force sufficient to remove its gas whereas H I-rich ones do not. This suggests a strong link between ram pressure stripping and satellite H I removal, which our results corroborate from a different perspective. We note that halos start to have very small amount of hot gas starting with halo masses of $\sim 10^{11} M_{\odot}$ from which some environmental effect might slightly be seen (which is the case on Fig. 5)

There are emerging observations of the relationships between H I content in centrals and satellites as a function of halo mass. In particular, Catinella et al. (2013) combined data from GASS⁷ with the Yang et al. (2007) SDSS halo catalogue to separate their galaxies into satellites and centrals. In Figure 6, we compare M_{HI} and M_{HI}/M_* as a function of M_{halo} for the combined sample of centrals and satellites to these data. The blue lines show the median values for all the simulated galaxies (with the error bars showing cosmic variance), while the red stars show the medians of all the observational data (with the uncertainties computed using jackknife resampling over 8 subsamples). To make a fairer comparison to the data, we restrict our sample to $M_* > 10^{10} M_{\odot}$ galaxies. In their sample, galaxies with $\log(M_*/M_{\odot}) \leq 10.5$ have a fixed H I mass lower limit value of $\log(M_{\text{HI}}/M_{\odot}) = 8.7$ whereas for more massive galaxies, the H I limit is derived from $\log(M_{\text{HI}}/M_*) = -1.8$; we clip our simulated values according to those limits and added a random scatter of 0.3, which we adopt to produce a comparable scatter to that seen in the lower limits within the GASS sample. In this way, the simulated galaxies having H I mass below the GASS detection limit are still included in the simulated sample, with their gas content increased to the observational detection limit. However, because we are considering medians, this does not have a major impact on the results, except in the highest mass bin where the median simulated galaxy is a non-detection according to the GASS limit. In this case, had we taken the actual M_{HI} value rather than the GASS limit, the medians of both M_{HI} and H I richness would be lowered by ~ 1 dex.

Figure 6 shows that the models qualitatively track the data across the halo mass range probed, but there are systematic discrepancies. For the observations, the uncertainty in the median is computed via jackknife resampling over 8 subsamples, which yields an extremely small statistical uncertainty. These do not include potential systematic errors such as cosmic variance and uncertainties in the measurements themselves, which would likely make the errors significantly larger. The model error bars show cosmic variance over the eight simulation sub-octants. There are also systematic uncertainties in computing the H I content from our models, as our radiative transfer approximation is necessarily crude and can introduce uncertainties potentially as large as several tenths of dex. Hence, the discrepancy in the predicted M_{HI} of no more than 0.2 dex, and that for H I richness peaking at 0.3 dex in the poor group ($\sim 10^{13} M_{\odot}$) regime, represents encouraging if not perfect agreement.

While there is general agreement, these discrepancies may prove enlightening. In the lower panel, we notice that the H I richness predicted in the simulation is somewhat above the observations (at most $\sim 2 \times$ H I richer) for halo masses around $M_{\text{halo}} < \sim 10^{13} M_{\odot}$, being the same in the lowest halo mass range. The decrease in gas fraction from the simulation seems to be lower than that in the observa-

⁹ <https://science.nrao.edu/science/surveys/littlethings>




Figure 6. A comparison of simulated galaxies with GASS observations from Catinella et al. (2013) for galaxies with $M_* \geq 10^{10} M_\odot$. We show the H I mass (top) and H I richness (bottom) as a function of halo mass. The non-detected sources and the corresponding simulated galaxies are included in the median. Blue solid lines are running medians for simulated galaxies while the red stars show median values for the GASS data. Error bars are obtained from cosmic variance (for simulation) and jackknife resampling (observations). The dashed black lines are running medians for simulated galaxies using a model without quenching (and a slightly different wind model).

tional data. In the high mass halos ($M_{\text{halo}} > 10^{13} M_\odot$), the observational data show a rather constant H I richness regardless of the halo mass, while that of the simulations decreases with halo mass owing to the growing population of H I poor galaxies as explained previously and also shown in Figure 4. However, this is partially driven by the fact that many of the observed galaxies in this bin are non-detection placed at their upper limit of $\log(M_{\text{HI}}/M_*) = -1.8$. While we have tried to mimic GASS selection as closely as possible, it is difficult to assess the robustness of a comparison driven by non-detections. However, it may suggest that the predicted satellite H I content in this crucial poor group mass range is somewhat discrepant with the data. At lower and higher halo masses, the simulation agrees better with the observations.

So far, we have included the undetected sources in the comparison mimicking the way GASS handled such cases. For most bins this had little effect, but at the massive end there were many non-detections. To investigate this, we present in Figure 7 a comparison between the simulations and the observations, now excluding sources that were undetected in H I and the corresponding galaxies in the simulations. We further separate the sample into central and satellite galaxies. The blue lines represent the simulations, with the diamonds representing the central galaxies and the circles the satellite galaxies. The red lines represent the ob-

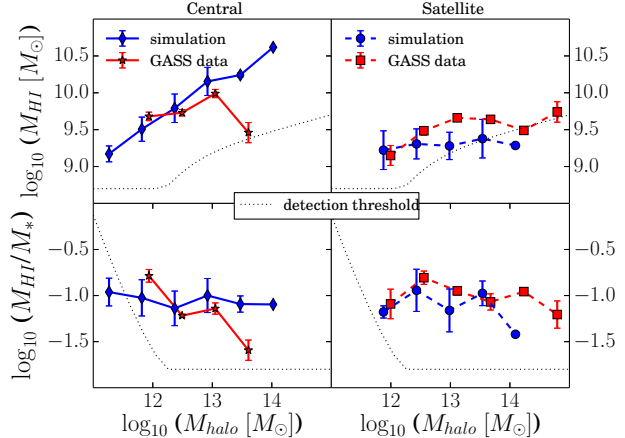


Figure 7. Similar to Figure 6 showing H I richness versus halo mass, here excluding galaxies below the observational detection limit, and also showing the breakdown between central and satellite galaxies. Blue lines represent the simulation (*diamonds*: centrals; *circles*: satellites) and red lines the GASS data (*stars*: centrals; *squares*: satellites). Considering the rather small stellar mass range and the limited H I mass, all of the medians show a flat relationship between H I-richness and halo mass. The simulations generally agree with the observations except for central galaxies within $\geq 10^{13} M_\odot$ halos, though there are modest systematic offsets in the satellite populations. The black dotted lines approximately shows the GASS detection threshold.

servational data, with the stars indicating the centrals and the squares the satellites. The error bars show 1σ uncertainties in the median within each bin, as described for Figure 6. We also show as the black dotted lines the GASS detection threshold, as mentioned previously. We assume that a given stellar mass corresponds to a halo mass as determined by Yang et al. (2008, see their eq. 7) from their SDSS halo catalog, namely:

$$M_{*,c} = M_0 \times \frac{(M_h/M_1)^{\alpha+\beta}}{(1 + M_h/M_1)^\beta} \quad (6)$$

with $[\log M_0, \log M_1, \alpha, \beta] = [10.86, 12.08, 0.22, 1.61]$. At high halo masses, the GASS data are mostly near the detection threshold.

The simulation and the observations are again in reasonable agreement (though not within the formal statistical uncertainties), with the exception of central galaxies in halos greater than $10^{13} M_\odot$. This plot at face value shows weak environmental effects on the satellite galaxies, as the median satellite has a roughly invariant H I mass and richness. However, this is only for the *detected* population. As discussed in Figure 4, the simulations predict that the main environmental effect on satellites consists of a fairly rapid removal of H I, which means that the galaxies that retain their H I are not substantially different. However, the *fraction* of galaxies with H I is a strong function of halo mass, as seen in Figure 4.

Returning to Figure 6, we now examine the impact of our quenching model on these results by comparing the vzw model without quenching to our standard ezw model. The vzw model (dashed lines) is slightly lower than our primary model in both gas content and H I richness. This is consistent with the results from Davé et al. (2013), who found

that stronger outflows at low masses, as assumed in our ezv model, yield greater H I masses in galaxies. In this case, the selection criterion (corresponding to $M_{\text{halo}} \gtrsim 10^{12} M_{\odot}$) emphasises galaxies that are most susceptible to being quenched. Despite the quenching in the ezv model, the H I content in this model is higher than in the vzw model, which likely stems from the lower mass loading factors present in the vzw model at low galaxy masses, which simply gets propagated to higher masses¹⁰. Hence, our quenching prescription does not seem to have a major effect on the H I content of high-mass galaxies. Coincidentally, the vzw model actually matches these data somewhat better, though we will show later that it does significantly worse in comparisons with other environmental measures.

In summary, the halo mass is a strong determinant for the H I content of its galaxies. For all the simulated galaxies in $M_{\text{halo}} \lesssim 10^{12} M_{\odot}$ halos, the H I mass in centrals rises with halo mass while the H I richness drops, and the H I richness of satellite galaxies is essentially unaffected by being within another halo. In more massive halos, the H I mass of central galaxies continues to rise but more slowly, and the median H I richness in satellites drops dramatically, driven by the bimodal appearance of substantial numbers of H I-poor satellites. This trend is likely driven primarily by gas stripping, as we discuss further below; our ad hoc quenching model has little impact on this. The predicted H I content of satellites in massive halos is in reasonable agreement with observations when matching the data selection, with some modest discrepancies particularly at high halo masses. The GASS galaxies show a less strong trend versus halo mass (Catinella et al. 2013), but when we apply similar detection limits to our simulations, the agreement is reasonable though not perfect. While these trends in massive halos have already been seen in observations, they provide a valuable test for models, and at low masses our simulations make predictions for the H I content as a function of halo mass and environment that could be tested against future surveys.

3.3 Timescales for H I loss in satellites

The bimodality of H I richness in satellites of high-mass halos indicates a fairly rapid loss of H I once a satellite falls into a $\gtrsim 10^{12} M_{\odot}$ halo. In this section, we trace galaxies back in time to determine the evolution of their H I content once they become satellite, to quantify the timescales for H I removal.

In Figure 8, we show the evolution of the H I richness of satellite galaxies since they were last central ($t = 0$). We only consider galaxies that were always central before $t = 0$ and that stay satellite until $z = 0$. Then, in later outputs, we compute the H I richness relative to the H I richness at $t = 0$ ¹¹. We plot this in Figure 8 for three different halo mass bins: $10^{11} M_{\odot} \leq M_{\text{halo}} < 10^{12} M_{\odot}$, $10^{12} M_{\odot} \leq M_{\text{halo}} < 10^{13} M_{\odot}$ and $10^{13} M_{\odot} \leq M_{\text{halo}}$, where the halo masses are those at the time of infall. The median H I richness as a function of time is indicated by the red (with




Figure 8. The evolution of H I richness and sSFR for satellite galaxies after infall into halos with masses $10^{11} M_{\odot} \leq M_{\text{halo}} < 10^{12} M_{\odot}$ (red with circles), $10^{12} M_{\odot} \leq M_{\text{halo}} < 10^{13} M_{\odot}$ (green with squares) and $M_{\text{halo}} \geq 10^{13} M_{\odot}$ (blue with downward triangles). Thick solid lines show median values for the H I richness relative to the value at $t = 0$, defined for each galaxy as the time when it becomes a satellite (left axis). The green vertical hatched area represents the 1σ dispersion for $10^{12} M_{\odot} \leq M_{\text{halo}} < 10^{13} M_{\odot}$. Dashed lines (with the colours and markers corresponding to the halo mass bins) correspond to median values for the sSFR relative to the value at $t = 0$ (right axis). The black cross hatched area represents the 1σ uncertainty of the sSFR in host halos of mass $\sim 10^{12} M_{\odot}$. The yellow line plots the H I richness fraction of central galaxies (oblique hatched area = 1σ scatter) that have never been satellite, with $t = 0$ the time when their stellar masses are comparable to those of satellites at infall into host halos of mass $\sim 10^{12} M_{\odot}$. This is used as a reference. Satellites take less time in massive halos to remove their gas.

circles), green (with squares), and blue (with triangles) lines for the increasing halo mass bins, respectively. The green vertical hatched region represents the 1σ variation around the median for the $10^{12} M_{\odot} \leq M_{\text{halo}} < 10^{13} M_{\odot}$ case; the others are similar. The yellow line (with the 1σ uncertainty shown by the oblique hatched area), for reference, shows the H I richness of central galaxies that have never been satellite, where $t = 0$ is the time when the galaxies' stellar masses are comparable to the satellites' stellar masses upon infall into a $\sim 10^{12} M_{\odot}$ halo.

Satellite galaxies in massive halos take less time to have their H I removed than lower halo mass galaxies. There is only a small difference between $10^{11} M_{\odot} \leq M_{\text{halo}} < 10^{12} M_{\odot}$ and $10^{12} M_{\odot} \leq M_{\text{halo}} < 10^{13} M_{\odot}$ halos, having a timescale of 1-2 Gyr to be lowered to half of their initial H I richness, while for $\geq 10^{13} M_{\odot}$ halos the timescale is significantly less than a Gyr. Hence, gas removal begins quite rapidly, often in less than a single halo dynamical time, but nonetheless full stripping of the H I (e.g. less than 10% of initial H I left) does not occur until after several Gyr. Indeed, in the lowest mass bin, even after 5 Gyr the H I is still typically $\gtrsim 10\%$ of the initial value. Note that the scatter about the median is quite large, so that the typical time to lose half of the H I at 1σ above the median is > 3 Gyr instead of ~ 1.5 Gyr. Hence, at least some satellites continue to retain significant H I for many Gyr even in high mass halos. It is also worth mentioning that there is a significant difference even between the satellite galaxies in the lowest halo mass bin and the central

¹⁰ A comparison between the ezv and vzw models as a function of stellar mass was shown in Davé et al. (2013).

¹¹ The scatter at $t = 0$ is a plotting artefact owing to the binning; all values at $t = 0$ are unity by definition.

galaxies having similar stellar mass (yellow region), showing that there is some environmental effect on satellites even in lower-mass halos. We note that some SAMs have adopted the prescription that hot gas is fully stripped immediately or within a halo dynamical time when it enters into a massive halo (e.g. De Lucia et al. 2012), but this is an oversimplification that is unlikely to be correct even in the median case. Recent SAMs have begun to include more sophisticated descriptions of gas stripping (e.g. Henriques et al. 2015). Finally, it is important to recall from Figure 4 that the distribution of H I richnesses in satellites is bimodal, meaning that the median evolution more broadly tracks the *fraction* of satellites that are having their H I stripped, rather than the evolution of any individual satellite. Hence, the timescales here should be viewed as the typical timescale after infall at which H I is stripped, but the stripping itself may happen more rapidly.

Since the decay is roughly log-linear, we can fit an exponential decay timescale for the H I richness of the form $R_{\text{HI}} = e^{-t/\tau}$, where R_{HI} is the ratio of M_{HI}/M_* relative to that at $t = 0$. For the H I richness, the e-folding decay timescales are 2.63, 2.01, 0.70 Gyr for the $10^{11} M_\odot \leq M_{\text{halo}} < 10^{12} M_\odot$, $10^{12} M_\odot \leq M_{\text{halo}} < 10^{13} M_\odot$ and $10^{13} M_\odot \leq M_{\text{halo}}$ bins, respectively, thus quantifying the more rapid H I loss in massive halos. In low mass halos, the timescales are comparable to or longer than the halo dynamical time of $\sim 2.5(1+z)^{-3/2}$ Gyr (Finlator & Davé 2008), indicating that gradual processes such as starvation are the drivers, while in more massive halos it is much shorter than a halo dynamical time, suggesting a more efficient and local process. When taking a fix stellar mass range $10^9 M_\odot \leq M_* \leq 10^{10} M_\odot$ at $z = 0$ (chosen to ensure that there are still many galaxies at the time of infall above the resolution limit), the e-folding H I richness decay timescales become 2.88, 1.84, 0.26 Gyr for the $10^{11} M_\odot \leq M_{\text{halo}} < 10^{12} M_\odot$, $10^{12} M_\odot \leq M_{\text{halo}} < 10^{13} M_\odot$ and $10^{13} M_\odot \leq M_{\text{halo}}$ bins. Particularly for galaxies with $10^9 M_\odot \leq M_* \leq 10^{10} M_\odot$ in massive halos (at $z = 0$), it takes a very short time for their H I to be stripped.

To better understand the physical origin of the decline in H I, we can compare the decay in H I richness to that of the star formation rate. If the decline in H I is much faster than the decline in SFR, then this indicates preferential removal of H I relative to star-forming gas, suggesting stripping. If the decline rates are comparable, then this can be most simply interpreted as starvation. In the absence of environmental effects, galaxies are in a steady state of accretion versus consumption (e.g. Dekel et al. 2009; Davé et al. 2012; Lilly et al. 2013). But if the accretion is truncated owing to environmental effects, leading to starvation, then as the H I runs out, the star-forming gas will deplete, and hence the SFR will commensurately drop.

In Figure 8, the dashed lines analogously show the evolution of the specific SFR ($\equiv \text{SFR}/M_*$) in satellites once they fall into halos (the colours show the respective halo mass bins). For the lower mass halos, the timescale for sSFR decay is virtually identical to that of H I richness decay (though we see a small discord between sSFR and H I richness for the $10^{12} - 10^{13} M_\odot$ halo mass bin after 2 Gyr). This indicates that these galaxies are simply consuming their gas and it is not being replenished, resulting in a coincident drop in gas and SFR. Hence in $M_{\text{halo}} \lesssim 10^{12} M_\odot$ halos, star-

vation appears to be the key mechanism for H I attenuation, which occurs because the streams that feed low-mass halos generally feed the central galaxy rather than satellites. This is likely the origin of the residual environmental effect that is evident from the difference in the satellites' and centrals' evolution (red or green vs. yellow line).

For high mass halos, in contrast, the H I attenuates much more quickly than the sSFR. The post-infall evolution of sSFR is essentially independent of halo mass¹², indicating that the denser gas and galaxy environment around more massive halos do not markedly affect gas consumption from an existing star-forming gas reservoir. The more rapid drop in H I suggests that ram pressure stripping is at work, since this should affect the loosely-bound H I much more than the dense star-forming gas. Hence the onset of a hot gaseous halo, even in fairly poor groups with $M_{\text{halo}} > 10^{13} M_\odot$, is already sufficiently to produce substantial ram pressure stripping (Cunnama et al. 2014) of infalling H I. Nonetheless, the SFR drop is mostly unaffected by ram pressure stripping, and remains consistent with simple starvation even in massive halos. The difference between the physics that results in the decline in SFR versus the decline in H I in massive halos is an important prediction of these simulations, and can perhaps reconcile conflicting results regarding the ubiquity of ram pressure stripping in H I cluster studies (e.g. van Gorkom et al. 2003) versus the much longer decay timescale of satellite star formation in massive halos (e.g. Wetzel et al. 2013).

3.4 H I as a function of environment

In section 3.2, we found that the H I content of galaxies is reduced at high halo masses. Here we examine a related quantity, which is the H I content as a function of environment. Since halo mass is correlated with the environment, e.g. as measured by the number of nearby galaxies, the reduction in H I at high halo masses is expected to be qualitatively mimicked in high density regions. Such a trend agrees with the observations of Fabello et al. (2012), which are based on stacked ALFALFA H I data. However, they inferred that environmental processes begin at $M_{\text{halo}} \sim 10^{13} M_\odot$, whereas our simulations predict that the suppression starts at $M_{\text{halo}} \sim 10^{12} M_\odot$. In this section we conduct a more detailed comparison to Fabello et al. (2012), to better understand what these data may be telling us about the impact of environment on H I content.

Figure 9 shows the median H I richness and specific SFR of galaxies as a function of their environment parameter N from Fabello et al. (2012), relative to the H I richness and sSFR of $N = 0$ galaxies. N is defined to be the number of galaxies with $M_* > 10^{9.5} M_\odot$ within a cylinder of radius 1 Mpc and a redshift path length of $\pm 500 \text{ km s}^{-1}$. In our simulation, the mean halo mass for $N \geq 7$ (including galaxies with stellar masses $10^{10} M_\odot \leq M_* \leq 10^{10.5} M_\odot$) is $\sim 10^{13.2} M_\odot$, which is consistent with Figure 8 in Fabello et al. (2012). We choose two galaxy stellar mass bins, $M_* = 10^{10} - 10^{10.5} M_\odot$ and $M_* = 10^{10.5} - 10^{11} M_\odot$ to compare to the observational data presented in Fabello et al.

¹² Figure 8 shows that the evolution of the sSFR ratios are all within 1σ uncertainty of $10^{12} M_\odot \leq M_{\text{halo}} < 10^{13} M_\odot$.




Figure 9. The median H I richness (blue) and sSFR (red) of galaxies as a function of environmental parameter N (see text) relative to their value for $N = 0$ galaxies. Solid lines (with circle markers) correspond to our simulated galaxies and stars to observational data from Fabello et al. (2012). The errorbars are 1σ uncertainties. Left and right panels correspond to galaxies with stellar masses of $10^{10} M_{\odot} < M_* < 10^{10.5} M_{\odot}$ and $10^{10.5} M_{\odot} < M_* < 10^{11} M_{\odot}$, respectively. The bottom panels only include galaxies with $sSFR \geq 10^{-11.2} \text{ yr}^{-1}$. The simulations show a good agreement with the observational data on the H I content fraction for star-forming galaxies but they show some excess in H I richness for high N in quenched halos.

(2012). In Figure 9, blue represents the H I richness, and red the sSFR, with the solid lines (with circle markers) showing results from the simulation and the stars those from the observations. The errorbars represent 1σ uncertainties. In the upper panels we apply no minimum sSFR selection criteria, while in the lower panels we select only the galaxies with $sSFR \geq 10^{-11.2} \text{ yr}^{-1}$, i.e. it excludes quenched galaxies. Note that the observational data determine the global sSFRs from SED fitting over the whole galaxy.

In both the simulations and data, there is an overall trend of dropping median H I richness and sSFR in denser environments. However, the simulation results are typically slightly higher than the observations particularly at high- N , implying that galaxies located in dense regions have somewhat too much H I. This discrepancy is essentially all being driven by the quenched galaxies; in the lower panels, we see that for star-forming galaxies, the predicted H I trends of the simulation are in very good agreement with the data, showing much less decline in dense regions.

The discrepancy in dense regions may be related to the excess in H I richness at $M_{\text{halo}} \sim 10^{13} M_{\odot}$ seen in Figure 6. This further indicates that, while overall the simulation broadly attenuates the H I content as observed, in detail there may be some discrepancies, especially in our quenched galaxies.

A comparison with our vzw model, which excludes quenching, in Figure 10, shows the effect of quenching for a range of environmental conditions. Without quenching




Figure 10. Similar to Figure 9, showing mean H I richness as a function of N relative to $N = 0$ galaxies, for two simulations: Solid lines with circles correspond to our fiducial model (ezw) and dashed lines with squares correspond to a model without quenching (vzw). We also show observational data from Fabello et al. (2012) as represented by the stars. The effect of quenching is qualitatively stronger at high N than at high halo mass (cf. Fig. 6)

(vzw), the change in H I richness relative to the H I richness at $N = 0$ towards denser regions is more pronounced (for $10.5 < \log_{10}(M_*/M_\odot) < 11$). This shows that the quenching model has a comparable qualitative impact with respect to environment as it does with respect to halo mass (cf. Fig. 6). For lower mass galaxies ($10 < \log_{10}(M_*/M_\odot) < 10.5$), there is very little difference between the quenching (ezw) and no-quenching (vzw) simulations, since the quenching mass scale we impose is at the upper end of this mass range. For more massive galaxies (right panels), however, the difference between the two is larger. The vzw model actually fares better here, suggesting again that the quenching model in ezw is not in good agreement with the H I observations. The ezw model does fare better, however, than the vzw model for the star-forming sample, which is not surprising as it produces a better match to the stellar mass function below L^* (Davé et al. 2013). Given the ad hoc nature of our quenching prescription, one should not read too much into these trends, but an important broader point is that observations of H I versus environment could potentially be used to constrain models of star formation quenching in central galaxies.

3.5 H I radial halo profiles

So far we have focused on the properties of the total H I content within galaxies and halos. To investigate the physical processes by which environment affects H I in more detail, here we study the radial distribution of H I within halos.

Figure 11 shows the H I richness of satellites as a function of the radial distance from the halo centre, in units of the halo’s virial radius, defined in §2.1 (similar to what was done in Solanes et al. 2001). The main trend shows that at all radii, the H I richness always decreases with increasing halo mass. This is the same overall trend as shown in Figure 4, except here we additionally show that it occurs as a function of radius.

Besides this overall reduction, there is a noticeable change in the H I richness profile with halo mass. In the low mass halos ($M_{\text{halo}} \lesssim 10^{12} M_\odot$), we see that the H I richness of satellite galaxies is independent of their distance to the centre of the halo, consistent with there being essentially no stripping of their gas content as the satellite falls deep into these halos. In contrast, in high mass halos there is a clear decrease in H I richness at small radii, with the trend gradually changing from no radial trend to a very steep one with halo mass. For the most massive halos ($M_{\text{halo}} \sim 10^{14} M_\odot$), satellite galaxies close to their centre are more than two orders of magnitude lower in H I richness than those near the virial radius, with the trend mildly steepening within the inner third of the virial radius. This decrease of H I fraction for the most massive halos may owe to the fact that galaxies closer to the centre of the halo accreted earlier than those at the outskirts, leading to a longer time for stripping to take effect.

We also show the results from vzw model in Figure 11 as dashed lines, to check the impact of quenching on the H I-richness profile. At all radii and for most of the halo mass range, vzw-galaxies have approximately half an order of magnitude lower H I-richness than ezw galaxies. Hence the radial trend does not seem to be much affected by our quenching prescription. This is not surprising because




Figure 11. The H I richness of satellites as a function of radial distance from the halo centre in units of the virial radius. Different colours show mean values for different halo mass bins, as indicated by the colourbar. The solid lines show our fiducial model (ezw) and the dashed lines show the model without quenching (vzw). The H I richness of satellite galaxies is independent of radius for smaller halos, but decreases by ~ 100 times near the centre of a massive halo ($M_{\text{halo}} \sim 10^{14} M_\odot$) compared to the outskirts.

our quenching prescription does not directly affect satellite galaxies since they are typically below our σ quenching threshold.

In summary, the radial profile of H I richness shows no trend for low-mass halos, but at $M_{\text{halo}} \gtrsim 10^{13} M_\odot$ there is a marked drop in H I richness towards the centre which becomes rapidly more pronounced at higher masses. This is qualitatively consistent with the idea that gas stripping is quite efficient in high mass halos that are abundant in hot hydrostatic gas. Further comparisons to observations will elucidate whether these simulations are capturing all the relevant processes accurately.

4 H I EVOLUTION AND MERGERS

4.1 Evolutionary tracks

Figure 12 shows evolutionary tracks from $z = 5 \rightarrow 0$ for selected individual simulated galaxies spanning a range of final stellar masses and halo masses. The galaxies are selected to show different features such as different final stellar masses, different halo masses, the effects of quenching and recent and early mergers. The thick black lines show the H I mass evolution, the thick grey lines show the stellar mass growth and the thin black lines show the SFR evolution (scale on the right axis). We indicate along the H I mass lines (thick black lines) whether the galaxy was central (red circles) or satellite (blue stars) at that epoch. Vertical green lines show where greater than 1 : 3 mergers occur, as described in §2.3. We also show the log of the halo and stellar masses at $z = 0$ by the two numbers in the top right corner, where the colour of the numbers indicates whether the galaxy is central (red) or satellite (blue) at that redshift. Note that there are two roughly Milky Way-sized galaxies in terms of halo and stellar mass, in panels (k) and (l).

The tracks show that galaxies grow in H I and stellar

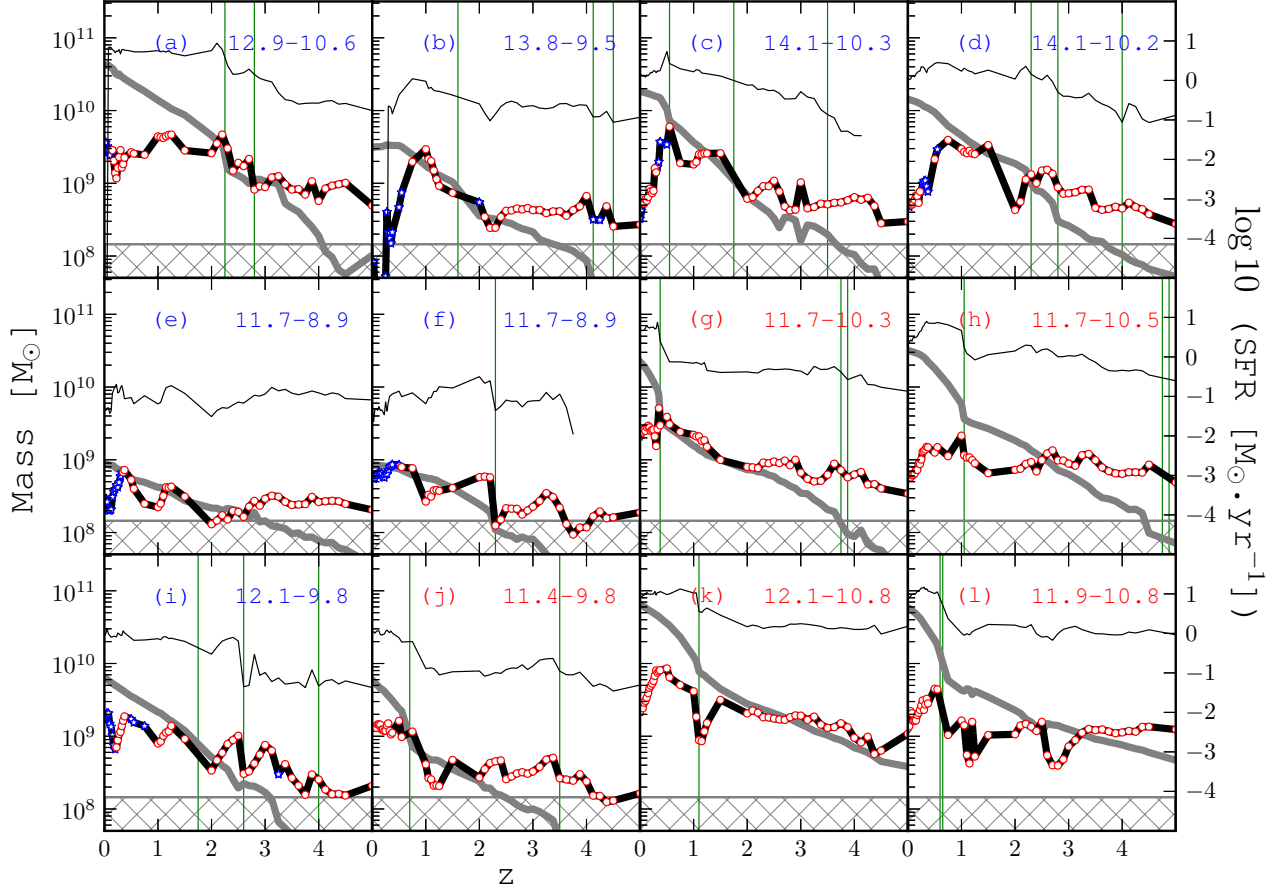


Figure 12. The growth of 12 individual galaxies from $z=5$ to $z=0$. The thick black lines show the H I mass evolution, the thick grey lines show the stellar mass evolution, and the thin black lines show the SFR (scale on the right axis). The green vertical lines indicate the redshifts at which galaxies have undergone a $> 1:3$ merger. The numbers on the top right corner are the halo mass and stellar mass at $z=0$, respectively (the colour indicates whether the galaxy is central (red) or satellite (blue) at $z=0$). Red circles along the H I mass lines represent that the galaxy is a central at the given redshift and blue stars show that the galaxy is a satellite.

mass together, although in general the stellar mass grows more quickly than the H I mass. We will show in the next section that in these simulations, this evolution is a result of the roughly constant relationship between H I and stellar mass, and is not driven by redshift evolution of this relation. That is, as the galaxy grows in M_* , it typically has lower M_{HI}/M_* , and thus its H I mass has grown more slowly than stellar mass. There is also a relationship between H I growth and SFR, and in particular a merger (identified by the vertical lines) sometimes drives a rise in the SFR but can result in either a rise or decline in H I mass; we will quantify this in §4.3. Another aspect depicted is the quenching prescription, as described in §2.1, which stops galaxy from forming stars when it is located in a massive halo. For instance, in panel (a), we can see that the galaxy is quenched just before $z = 0$ once it entered a massive halo, but its H I gas is still available to possibly feed star formation before complete stripping occurs. Conversely, in panel (b), the SFR starts decreasing in concert with the H I, and then eventually is fully truncated as the H I is removed. The galaxy ceases to form stars and the stellar mass even decreases as some of the stars get stripped. Stripping is also seen in panels (c) & (d). In these two cases, the SFR does not go to zero, but we

see that once the galaxy crosses into a bigger halo the gas is quickly exhausted while the SFR decreases at a far lower rate: this confirms what we showed in Figure 8 with the blue solid and the blue dashed lines for the most massive halos.

For central galaxies becoming satellites by entering into low-mass ($\lesssim 10^{12} M_\odot$) halos, Figure 8 showed that the decrease in gas fraction is very much similar to the decrease of specific SFR. To illustrate this, two central galaxies that have become satellites in less than $10^{12} M_\odot$ host halos are shown in panels (e) & (f): the H I content decreases along with the SFR. We speculate that these galaxies do not receive any gas infall, and hence the decrease in their gas content owes to the fact that they are still producing stars without much further gas replenishment. Galaxies like (i) show highly fluctuating gas content as well as star formation rate, owing to mergers or fly-by's where some of the particles of either of the galaxies get disrupted or even stripped from the galaxies. The two Milky Way-like galaxies (k) and (l) both undergo mergers at fairly late epochs, but in (l) the post-merger galaxy then reduces its H I mass, whereas in (k) it gradually increases after the merger. The galaxy in panel (j) undergoes a merger at $z = 0.7$ but, at least down to $z = 0$, the H I richness seems unaffected by this.

These tracks illustrate the diversity in trends that impact the H I evolution in galaxies that we explore next, and in particular the relationship between mergers and H I that we quantify below.

4.2 H I richness evolution with redshift

Our results above suggest that massive halos with hot gas are effective at stripping the H I from infalling satellites, while $M_{\text{halo}} \lesssim 10^{12} M_{\odot}$ starve their satellites of H I over longer timescales. This explains the trend in §3.2 where the H I richness drops significantly for satellites relative to centrals in halos above this mass. In this section, we examine at which cosmic epoch this trend appears, and how it evolves to the present day.

Figure 13 shows the evolution of the median H I richness versus halo mass (only including galaxies with $M_{*} \geq 1.45 \times 10^8 M_{\odot}$ at the respective redshift), similar to that in the lower panel of Figure 4 but now showing the redshift evolution from $z = 3$ to $z = 0$. In this plot, the stellar mass and halo mass of the galaxy are measured at the respective redshift. We take all of the galaxies regardless of whether they are progenitors of those at $z=0$ or not as long as $M_{*} \geq 1.45 \times 10^8 M_{\odot}$. As discussed in Davé et al. (2013), the H I richness at a given mass hardly evolves in time; in Davé et al. (2013) this was shown versus stellar mass, while here we show it versus halo mass, but as expected from abundance matching models (e.g. Behroozi et al. 2013) there is not much evolution in the $M_{*} - M_{\text{halo}}$ relation. At all redshifts, we see the trend of lower H I-richness in more massive systems, which continues roughly unabated over the mass range probed.

In contrast to centrals, satellites at all redshifts show a reduction in their H I richness relative to centrals at $M_{\text{halo}} > 10^{12} M_{\odot}$. This is likely tied to the emergence of hot gaseous halos at this mass scale, which is roughly independent of redshift (Kereš et al. 2005). At the highest redshifts probed here ($z = 3$), we do not have sufficient volume to probe satellites in halos significantly above this mass, but it seems that the trend remains consistent. We note that Davé et al. (2013) demonstrated that the H I-poor fraction of satellites occurs regardless of the quenching prescription (which is based on σ , and so lower- σ satellites are generally unaffected); hence, this is not what is driving the decrease. Rather, the processes described in the previous section of starvation and ram pressure stripping are predominantly responsible.

Figure 14 shows the median H I richness of central and satellite galaxies from $z = 5 - 0$. We separate the galaxies into bins of $z = 0$ stellar mass as shown in the figure by the colour coded lines, and compute the medians for all those galaxies' main progenitors at each redshift. Hence, each line defines a fixed population of galaxies, although their environment varies over time. Here we expand our simulated sample at high- z to masses below our nominal resolution limit to extend the curves to high redshifts. If we were to limit our progenitors to only resolved galaxies the curves would not show a different trend, but would only be noisier.

The predominant trend is that galaxies are more H I-rich at early times. However, at $z \gtrsim 1.5$, the difference between satellites and centrals is very small, and the two track each other very well. This could be because many of those satel-




Figure 13. The median H I richness as a function of halo mass at different redshifts, from $z=3$ (orange) to $z=0$ (black). Solid and dashed lines correspond to central and satellite galaxies, respectively. Different populations of galaxies at different redshifts show no redshift evolution of H I richness with respect to host halo mass.

lites were centrals in their own halo at early times. To check this, we also show as the green line the median H I richness for only the galaxies that are satellites at each redshift, and that are above our M_{*} resolution limit. Note that this is not a fixed population, as there are fewer galaxies that are still satellites at higher redshifts, and hence the scatter gets larger at higher redshifts. Nonetheless, the trend of increasing H I richness is generally present at roughly the same rate for galaxies that are satellites at every epoch. This indicates that even though some satellites may be centrals at earlier epochs, in fact the satellites themselves are also more H I rich at higher redshifts.

At $z \lesssim 1.5$, the satellite galaxies' H I richness begins to depart significantly from that of the centrals in halos of all masses, as large potential wells grow that can more effectively strip satellites. Observational (Muzzin et al. 2014) and theoretical (e.g. Bahé & McCarthy 2015) studies also suggest that gas stripping can be effective in cluster mass halos starting at these early epochs. An interesting feature we see is the flattening in H I richness evolution at $z \lesssim 1$ for the lowest mass central galaxies; the origin of this is not entirely clear. The lowest mass satellites, meanwhile, show the strongest drop to $z = 0$, as they are most susceptible to stripping processes.

We can also examine the evolution in bins of halo mass rather than stellar mass, as shown in Figure 15. We focus here on low-mass satellites since they are most strongly impacted by H I removal processes, namely satellites with $10^9 M_{\odot} < M_{*} < 10^{10} M_{\odot}$. We choose this mass range because, while we resolve to lower masses at $z = 0$, we cannot track their history very far back in time before they fall below our resolution limit. Figure 15 plots the H I richness evolution from $z = 5 \rightarrow 0$ for these satellites in various $z = 0$ halo mass bins as indicated by the colour-coding. Again, we see that at $z \lesssim 1$ there is a marked drop in the H I richness of satellites, but only in the most massive halos ($M_{\text{halo}} \gtrsim 10^{13} M_{\odot}$). The redshift at which H I removal becomes important is a strong function of halo mass, con-




Figure 14. The median H I richness as a function of redshift for central galaxies (solid lines) and satellites (dashed lines) in different stellar mass bins defined at $z = 0$, increasing from $M_* \sim 10^{9.3} M_\odot$ (black) to $M_* \sim 10^{11.3} M_\odot$ (orange). The green line shows the median H I richness of all galaxies that are still satellites at that redshift, with $M_* \geq 1.45 \times 10^8 M_\odot$.




Figure 15. The median H I richness as a function of redshift for satellites in different $z = 0$ halo mass bins, increasing from $M_{\text{halo}} \sim 10^{11.5} M_\odot$ (black) to $M_{\text{halo}} \sim 10^{14} M_\odot$ (orange). We only include satellites in the stellar mass range $10^9 M_\odot < M_* < 10^{10} M_\odot$ at $z=0$ and their most massive progenitors at higher redshifts.

sistent with the idea that stripping begins when the halo mass crosses a mass threshold (Figure 13) coincident with the onset of a substantial hot gaseous halo.

4.3 H I in mergers

In this section we investigate the impact that major mergers have on the H I content of galaxies, and in particular we quantify how much scatter in the relationship between M_{HI} (or H I richness) and M_* is induced by mergers.

Figure 12 shows that mergers in many cases result in a burst of star formation. What this means for the H I is unclear: we might expect a decrease of H I content owing to gas consumption that is too rapid for replenishment, but it could be that the H I might be increased if this gas is fuelling the




Figure 16. The effects of major mergers on the H I content of galaxies. We compare the distribution of the ratio of H I richness after and before every output for all galaxies (black histogram) to the distribution of ratios after and before every $> 1 : 3$ merger (red histogram). Shaded regions contain 90% of the data for the control (forward slashed) and merger (backward slashed) samples, with spreads of ~ 0.44 dex and ~ 0.75 dex, respectively. Median values are indicated by the vertical dashed lines. Mergers increase the median H I richness of galaxies, but the most dominant feature is the increase in the width of the distribution interpreted as simple stochasticity.

burst. For instance, in Figure 12 (c) (at $z \sim 0.5$), (f) or (h) we can see variations in the decrease in gas content. Conversely to those cases, we also encounter situations where the gas content keeps increasing, such as in Figure 12(c) (at $z \sim 1.75$), (i) (at $z \sim 1.75$) or (k). One possible impact for mergers may, therefore, be to increase the scatter in the relationship between H I and stellar mass. Here we more quantitatively consider how mergers alter the H I content of the galaxies.

Figure 16 quantifies the impact of mergers on the scatter in H I richness. Here we show the distribution of the ratios of the H I richness after and before every output (black histogram) compared to those after and before a greater than $1 : 3$ merger (red histogram). The distribution for all galaxies shows a tiny shift of the histogram towards values below 1 (median ≈ 0.95), which means that, on average, galaxies have their H I richness decreased slightly with time. This was anticipated if we refer to Figure 14 and Figure 15 where the H I richness decreases towards low redshift and serves as a reference for the distribution of the H I richness ratio after and before mergers.

The post-merger ratio (red histogram) shows two effects. First, the H I richness increases, with a median ratio of 1.21. This suggests that mergers generally boost the H I fraction, which occurs concurrently with an increase in SFR and a decrease in metallicity (Davé et al. 2013). This is consistent with the trend that high-H I galaxies have recently received a substantial input of new gas (Moran et al. 2012), and that such an input also fuels new star formation. It is less consistent with the idea that H I gets “used up” in a merger, even though this situation can be seen in some cases. The overall trend of higher H I content in higher SFR galaxies is consistent with observations (Robertson et al. 2013).

A second trend from Figure 16 is that the scatter in

relative H I richness is typically larger in post-merger galaxies than in the overall sample. This suggests that mergers do indeed tend to increase the scatter in H I richness at a given mass. The range of ratios enclosing 90% of galaxies rises from ~ 0.44 dex in the overall case to ~ 0.75 dex in the post-merger case.

A useful plot for understanding how H I participates in the baryon cycle is via a *deviation plot*, which quantifies second parameter variations in relationships versus a single quantity, such as stellar mass. An example of such a second parameter trend is the so-called fundamental metallicity relation (Mannucci & Cresci 2012; Lara-López et al. 2010), in which galaxies at a given stellar mass show lower metallicity at higher star formation rates.

Figure 17 shows two deviation plots. This quantifies the deviation of each galaxy from the overall trend given by the median value in each stellar mass bin. We first fit a cubic (third-order) spline to the median relationship between SFR and M_* , and M_{HI} and M_* . For each value of M_* , we then subtract the quantity from its corresponding median spline-fit value. This gives the deviation for each quantity, which we can then plot against each other.

The upper panel shows the deviation of the sSFR from its median value at a given M_* ($\Delta \log \text{sSFR}$) against the deviation of the H I richness from its median value at a given M_* ($\Delta \log(M_{\text{HI}}/M_*)$). The lower panel similarly shows a deviation plot of metallicity versus H I richness. To avoid congested plots, we chose to make density plots with two contours containing 1σ (black) and 2σ (blue) of the resolved galaxies. We fit a power law for all the galaxies (grey line), and for those which have just undergone a major merger between the last two outputs (red line).

From the top panel, we can see that recent mergers generally follow a different relation than the overall galaxy population. In particular, galaxies with a high sSFR also tend to have a high H I richness, but after a merger the sSFR is enhanced *more* than the H I richness. The grey line is similar to what was shown in Davé et al. (2013), which is explained that recent gas accretion tends to boost SFR. The point of Figure 17 is that mergers boost the SFR above the overall relation owing to general stochastic fluctuations in accretion, more so than the H I.

For comparison, the bottom panel shows a deviation plot of H I richness versus metallicity. This shows that galaxies with increased H I richness typically show a decreased metallicity at a given stellar mass. In contrast to the sSFR, here the trend for all galaxies is identical to recently merged galaxies, showing that low metallicity goes hand-in-hand with higher H I content independent of whether a merger happened or not. This is consistent with the idea that the recent infall that boosts the H I content also brings in lower metallicity gas. This shows that H I traces metallicity in the baryon cycle, while the SFR is specifically boosted by interactions.

Overall, mergers have a minor but noticeable impact on the H I content of galaxies, increasing it immediately after the merger, and adding to the scatter in H I richness. However, the increase in H I is essentially consistent with simple stochasticity in accretion within the baryon cycle, without there being an *additional* increase owing to the merger as seen in the case of the SFR. We caution that owing to the limited time resolution of our outputs, and possibly the low




Figure 17. The deviations of sSFR (top) and metallicity (bottom) from their median values at a given M_* against the deviation of the H I richness from its median value at a given M_* (see text for details). The density of galaxies in each deviation plot is indicated by the background colour scale, with the black and blue contours containing $\sim 68\%$ and $\sim 95\%$ of the galaxies, respectively. Yellow crosses and red circles show deviations for galaxy mergers with mass ratios $> 1 : 3$ and $> 1 : 1.33$, respectively. The grey and red solid lines show linear fits for all galaxies and major mergers, respectively. Merged galaxies sit just above the fitted line for all galaxies for the deviation in sSFR.

spatial resolution that is unable to fully resolve the galaxies' internal structure, this analysis may be underestimating the impact of mergers. More detailed simulations, and associated comparisons with data, can help us disentangle the roles of all the components of galaxies in mergers.

5 SUMMARY

We investigate the properties of H I in galaxies drawn from a cosmological hydrodynamic simulation including galactic outflows that reproduces many key observed properties of the H I. We focus on studying the relationship between H I and environment, quantified either via halo mass, central versus satellite galaxies, or local galaxy density. We also examine evolutionary trends in H I content versus redshift, and study how mergers influence the H I content. Our main results are summarised as follows:

- Our simulation shows very good agreement with the observed ALFALFA galaxy H I mass function (HIMF) and, furthermore, shows good agreement with the observed HIMF broken into stellar mass bins based on the GASS data. This suggests that the scatter between M_{HI} and M_* is generally well-reproduced in our simulation, providing a second-order test of our models.
- The H I content of central galaxies is governed primarily by the halo mass, with a positive correlation that is steeper

and tighter for $M_{\text{halo}} < 10^{12} M_{\odot}$. The median H I richness of centrals decreases slowly with halo mass, which continues up to the highest masses probed.

- For satellites, the medians of H I and H I richness show significantly lower values above $M_{\text{halo}} \sim 10^{13} M_{\odot}$, with the emergence of a bimodal distribution in which the drop in the median is driven by an increasing fraction of satellites essentially devoid of H I.

- When a galaxy falls into a more massive halo, its H I content is reduced. This trend occurs at all halo masses. However, the trend is strongly accelerated in halos with $M_{\text{halo}} > 10^{12} M_{\odot}$. The median e -folding timescale for removal is ~ 2 Gyr in lower mass halos, but only 0.7 Gyr in $M_{\text{halo}} \approx 10^{13} M_{\odot}$ halos.

- A comparison to star formation rate attenuation shows that the SFR is also attenuated once a galaxy becomes a satellite, but that there is no strong acceleration of this in massive halos. The physical implication is that star formation and H I attenuation are consistent with starvation in lower mass halos, while in more massive halos, the SFR still drops owing primarily to starvation while the H I is strongly affected by additional gas stripping processes associated with the presence of a hot gaseous halo.

- The H I richness at a given halo mass does not evolve much with redshift out to $z \sim 3$. There continues to be a rapid drop in a satellites' H I richness at $M_{\text{halo}} > 10^{12} M_{\odot}$. The median H I richness in all galaxies drops from $z = 5 \rightarrow 0$, with the satellites tightly tracking the centrals down to $z \sim 1$ and then dropping rapidly below that. The drop is tightly correlated with halo mass, starting first in the highest mass halos at $z \gtrsim 1$, while low-mass halos show no strong drop in the satellite H I richness at late times.

- In low mass halos ($M_{\text{halo}} \lesssim 10^{12} M_{\odot}$), the H I richness of satellite galaxies is independent of their distance to the central galaxy; whereas at high halo mass ($M_{\text{halo}} \gtrsim 10^{13} M_{\odot}$), galaxies become H I poorer towards the centre, indicative of strong gas stripping processes.

- Mergers typically cause a modest increase in the H I richness, while also increasing the scatter in H I richness. These deviations are consistent with mergers being an extreme example of a stochastic fluctuation in accretion, rather than being driven by internal processes particularly associated with the merger, as is the case with the SFR.

These results show that environment has a major impact on the H I content of galaxies. However, this impact is mostly confined to galaxies with $M_{\text{halo}} > 10^{12} M_{\odot}$. The fact that hot gaseous halos tend to develop around this mass scale in our simulations (Gabor & Davé 2012), as expected from virial shock cooling timescale arguments (Birnbom & Dekel 2003), suggests that the presence of hot gas is the main driver for environmental effects, particularly for satellite galaxies. The most straightforward interpretation of this is that ram pressure stripping effects are driving much of the H I satellite evolution in massive halos. But the timescale and manner in which the gas gets stripped, and its relationship to the molecular gas and stars, remains less clear. In future work we will examine gas stripping processes in more detail, using more physically-motivated models for star formation quenching and with simulations that use improved numerical techniques to better model the interaction of H I and hot ambient gas. This work sets the

baseline for such future studies, and highlights a new set of observational comparisons that can provide crucial constraints on the evolution of gas in galaxies.

ACKNOWLEDGEMENTS

The authors thank the referee for providing judicious comments and suggestions for additional work that helped improve this paper. The authors thank B. Catinella for providing us with GASS data for comparison and for helpful comments on the draft, and more broadly the GASS survey team for spurring this work with useful discussions. M. Rafieferantsoa acknowledges financial support from the South African Square Kilometer Array. This work was supported by the South African Research Chairs Initiative and the South African National Research Foundation. This work was supported by the National Science Foundation under grant number AST-0847667, and NASA grant NNX12AH86G. Computing resources were partially obtained through a grant from the Ahmanson Foundation.

REFERENCES

- Bahé Y. M., McCarthy I. G., 2015, *MNRAS*, 447, 969
 Behroozi P. S., Wechsler R. H., Conroy C., 2013, *ApJ*, 770, 57
 Berlind A. A., Weinberg D. H., Benson A. J., Baugh C. M., Cole S., Davé R., Frenk C. S., Jenkins A., et al., 2003, *ApJ*, 593, 1
 Birnboim Y., Dekel A., 2003, *MNRAS*, 345, 349
 Bower R. G., Benson A. J., Malbon R., Helly J. C., Frenk C. S., Baugh C. M., Cole S., Lacey C. G., 2006, *MNRAS*, 370, 645
 Catinella B., Schiminovich D., Cortese L., Fabello S., Hummels C. B., Moran S. M., Lemonias J. J., Cooper A. P., Wu R., Heckman T. M., Wang J., 2013, *MNRAS*, 436, 34
 Catinella B., Schiminovich D., Kauffmann G., Fabello S., Wang J., Hummels C., Lemonias J., Moran S. M., et al., 2010, *MNRAS*, 403, 683
 Chabrier G., 2003, *PASP*, 115, 763
 Cortese L., Catinella B., Boissier S., Boselli A., Heinis S., 2011, *MNRAS*, 415, 1797
 Cunnama D., Andrianomena S., Cress C. M., Faltenbacher A., Gibson B. K., Theuns T., 2014, *MNRAS*, 438, 2530
 Davé R., Finlator K., Oppenheimer B. D., 2011a, *MNRAS*, 416, 1354
 —, 2012, *MNRAS*, 421, 98
 Davé R., Katz N., Oppenheimer B. D., Kollmeier J. A., Weinberg D. H., 2013, *MNRAS*, 434, 2645
 Davé R., Oppenheimer B. D., Finlator K., 2011b, *MNRAS*, 415, 11
 Davé R., Oppenheimer B. D., Katz N., Kollmeier J. A., Weinberg D. H., 2010, *MNRAS*, 408, 2051
 De Lucia G., Blaizot J., 2007, *MNRAS*, 375, 2
 De Lucia G., Weinmann S., Poggianti B. M., Aragón-Salamanca A., Zaritsky D., 2012, *MNRAS*, 423, 1277
 Dekel A., Birnboim Y., Engel G., Freundlich J., Goerdt T., Mumcuoglu M., Neistein E., Pichon C., et al., 2009, *Nature*, 457, 451
 Dekel A., Silk J., 1986, *ApJ*, 303, 39

- Duffy A. R., Kay S. T., Battye R. A., Booth C. M., Dalla Vecchia C., Schaye J., 2012, *MNRAS*, 420, 2799
- Fabello S., Kauffmann G., Catinella B., Li C., Giovanelli R., Haynes M. P., 2012, *MNRAS*, 427, 2841
- Faucher-Giguère C.-A., Kereš D., Dijkstra M., Hernquist L., Zaldarriaga M., 2010, *ApJ*, 725, 633
- Fernández X., van Gorkom J. H., Hess K. M., Pisano D. J., Kreckel K., Momjian E., Popping A., Oosterloo T., et al., 2013, *ApJ*, 770, L29
- Finlator K., Davé R., 2008, *MNRAS*, 385, 2181
- Ford A. B., Oppenheimer B. D., Davé R., Katz N., Kollmeier J. A., Weinberg D. H., 2013, *MNRAS*, 432, 89
- Gabor J. M., Davé R., 2012, *MNRAS*, 427, 1816
- , 2015, *MNRAS*, 447, 374
- Giovanelli R., Haynes M. P., 1985, *ApJ*, 292, 404
- Giovanelli R., Haynes M. P., Kent B. R., Perillat P., Catinella B., Hoffman G. L., Momjian E., Rosenberg J. L., et al., 2005, *AJ*, 130, 2613
- Gunn J. E., Gott III J. R., 1972, *ApJ*, 176, 1
- Haardt F., Madau P., 2001a, in *Clusters of Galaxies and the High Redshift Universe Observed in X-rays*, Neumann D. M., Tran J. T. V., eds., p. 64
- , 2001b, in *Clusters of Galaxies and the High Redshift Universe Observed in X-rays*, Neumann D. M., Tran J. T. V., eds.
- Haynes M. P., Giovanelli R., Martin A. M., Hess K. M., Saintonge A., Adams E. A. K., Hallenbeck G., Hoffman G. L., et al., 2011, *AJ*, 142, 170
- Henriques B. M. B., White S. D. M., Thomas P. A., Angulo R., Guo Q., Lemson G., Springel V., Overzier R., 2015, *MNRAS*, 451, 2663
- Hess K. M., Wilcots E. M., 2013, *AJ*, 146, 124
- Hinshaw G., Larson D., Komatsu E., Spergel D. N., Bennett C. L., Dunkley J., Nolte M. R., Halpern M., et al., 2013, *ApJS*, 208, 19
- Hinshaw G., Weiland J. L., Hill R. S., Odegard N., Larson D., Bennett C. L., Dunkley J., Gold B., et al., 2009, *ApJS*, 180, 225
- Hopkins A. M., McClure-Griffiths N. M., Gaensler B. M., 2008, *ApJ*, 682, L13
- Hopkins P. F., Kereš D., Oñorbe J., Faucher-Giguère C.-A., Quataert E., Murray N., Bullock J. S., 2014, *MNRAS*, 445, 581
- Hopkins P. F., Quataert E., Murray N., 2012, *MNRAS*, 421, 3522
- Huang S., Haynes M. P., Giovanelli R., Brinchmann J., 2012, *ApJ*, 756, 113
- Hughes T. M., Cortese L., Boselli A., Gavazzi G., Davies J. I., 2013, *A&A*, 550, A115
- Katz N., Weinberg D. H., Hernquist L., 1996, *ApJS*, 105, 19
- Kennicutt R. C., Evans N. J., 2012, *ARA&A*, 50, 531
- Kereš D., Katz N., Weinberg D. H., Davé R., 2005, *MNRAS*, 363, 2
- Krumholz M. R., Gnedin N. Y., 2011, *ApJ*, 729, 36
- Lagos C. D. P., Baugh C. M., Lacey C. G., Benson A. J., Kim H.-S., Power C., 2011, *MNRAS*, 418, 1649
- Lagos C. d. P., Davis T. A., Lacey C. G., Zwaan M. A., Baugh C. M., Gonzalez-Perez V., Padilla N. D., 2014, *ArXiv e-prints*:1405.0016
- Lara-López M. A., Bongiovanni A., Cepa J., Pérez García A. M., Sánchez-Portal M., Castañeda H. O., Fernández Lorenzo M., Pović M., 2010, *A&A*, 519, A31
- Leitner S. N., Kravtsov A. V., 2011, *ApJ*, 734, 48
- Lemonias J. J., Schiminovich D., Catinella B., Heckman T. M., Moran S. M., 2013, *ArXiv e-prints*:1308.2178
- Leroy A. K., Walter F., Brinks E., Bigiel F., de Blok W. J. G., Madore B., Thornley M. D., 2008, *AJ*, 136, 2782
- Li C., Kauffmann G., Heckman T. M., Jing Y. P., White S. D. M., 2008, *MNRAS*, 385, 1903
- Lilly S. J., Carollo C. M., Pipino A., Renzini A., Peng Y., 2013, *ApJ*, 772, 119
- Mannucci F., Cresci G., 2012, *Memorie della Societa Astronomica Italiana Supplementi*, 19, 214
- Marcolini A., Brighenti F., D’Ercole A., 2003, *MNRAS*, 345, 1329
- Martin A. M., Papastergis E., Giovanelli R., Haynes M. P., Springob C. M., Stierwalt S., 2010, *ApJ*, 723, 1359
- Martin C., Engelbracht C., Gordon K., 2005, *Spitzer Proposal*, 20528
- McCarthy I. G., Frenk C. S., Font A. S., Lacey C. G., Bower R. G., Mitchell N. L., Balogh M. L., Theuns T., 2008, *MNRAS*, 383, 593
- Meyer M. J., Zwaan M. A., Webster R. L., Staveley-Smith L., Ryan-Weber E., Drinkwater M. J., Barnes D. G., Howlett M., et al., 2004, *MNRAS*, 350, 1195
- Mihos J. C., Hernquist L., 1996, *ApJ*, 464, 641
- Moore B., Katz N., Lake G., Dressler A., Oemler A., 1996, *Nature*, 379, 613
- Moran S. M., Heckman T. M., Kauffmann G., Davé R., Catinella B., Brinchmann J., Wang J., Schiminovich D., et al., 2012, *ApJ*, 745, 66
- Muratov A. L., Keres D., Faucher-Giguere C.-A., Hopkins P. F., Quataert E., Murray N., 2015, *ArXiv e-prints*:1501.03155
- Murray N., Quataert E., Thompson T. A., 2010, *ApJ*, 709, 191
- Muzzin A., van der Burg R. F. J., McGee S. L., Balogh M., Franx M., Hoekstra H., Hudson M. J., Noble A., Taranu D. S., Webb T., Wilson G., Yee H. K. C., 2014, *ApJ*, 796, 65
- Noeske K. G., Weiner B. J., Faber S. M., Papovich C., Koo D. C., Somerville R. S., Bundy K., Conselice C. J., et al., 2007, *ApJ*, 660, L43
- Obreschkow D., Croton D., De Lucia G., Khochfar S., Rawlings S., 2009, *ApJ*, 698, 1467
- Oppenheimer B. D., 2008, *PhD thesis*, The University of Arizona
- Oppenheimer B. D., Davé R., 2006, *MNRAS*, 373, 1265
- , 2008, *MNRAS*, 387, 577
- Oppenheimer B. D., Davé R., Katz N., Kollmeier J. A., Weinberg D. H., 2012, *MNRAS*, 420, 829
- Pappalardo C., Bianchi S., Corbelli E., Giovanardi C., Hunt L., Bendo G. J., Boselli A., Cortese L., et al., 2012, *A&A*, 545, A75
- Planck Collaboration, Ade P. A. R., Aghanim N., Armitage-Caplan C., Arnaud M., Ashdown M., Atrio-Barandela F., Aumont J., Baccigalupi C., Banday A. J., et al., 2014, *A&A*, 571, A16
- Popping A., Davé R., Braun R., Oppenheimer B. D., 2009, *A&A*, 504, 15
- Popping G., Somerville R. S., Trager S. C., 2014, *MNRAS*, 442, 2398

- Robertson P., Shields G. A., Davé R., Blanc G. A., Wright A., 2013, *ApJ*, 773, 4
- Rupen M. P., LITTLE THINGS Team, 2010, in *Bulletin of the American Astronomical Society*, Vol. 42, American Astronomical Society Meeting Abstracts #215, p. 202.04
- Salim S., Rich R. M., Charlot S., Brinchmann J., Johnson B. D., Schiminovich D., Seibert M., Mallery R., et al., 2007, *ApJS*, 173, 267
- Sancisi R., Fraternali F., Oosterloo T., van der Hulst T., 2008, *A&A Rev.*, 15, 189
- Schaye J., Dalla Vecchia C., Booth C. M., Wiersma R. P. C., Theuns T., Haas M. R., Bertone S., Duffy A. R., et al., 2010, *MNRAS*, 402, 1536
- Solanes J. M., Manrique A., García-Gómez C., González-Casado G., Giovanelli R., Haynes M. P., 2001, *ApJ*, 548, 97
- Springel V., 2005, *MNRAS*, 364, 1105
- Springel V., Hernquist L., 2003, *MNRAS*, 339, 289
- Tacconi L. J., Neri R., Genzel R., Combes F., Bolatto A., Cooper M. C., Wuyts S., Bournaud F., et al., 2013, *ApJ*, 768, 74
- van Gorkom J. H., Bravo-Alfaro H., Dwarakanath K. S., Guhathakurta P., Poggianti B. M., Schiminovich D., Valuri M., Verheijen M., et al., 2003, *Ap&SS*, 285, 219
- Vogelsberger M., Genel S., Springel V., Torrey P., Sijacki D., Xu D., Snyder G. F., Nelson D., et al., 2014, *ArXiv e-prints*:1405.2921
- Walter F., Brinks E., de Blok W. J. G., Bigiel F., Kennicutt Jr. R. C., Thornley M. D., Leroy A., 2008, *AJ*, 136, 2563
- Wang J., Fu J., Aumer M., Kauffmann G., Józsa G. I. G., Serra P., Huang M.-l., Brinchmann J., van der Hulst T., Bigiel F., 2014, *MNRAS*, 441, 2159
- Weiner B. J., Coil A. L., Prochaska J. X., Newman J. A., Cooper M. C., Bundy K., Conselice C. J., Dutton A. A., et al., 2009, *ApJ*, 692, 187
- Wetzel A. R., Tinker J. L., Conroy C., van den Bosch F. C., 2013, *MNRAS*, 432, 336
- Wiersma R. P. C., Schaye J., Theuns T., Dalla Vecchia C., Tornatore L., 2009, *MNRAS*, 399, 574
- Yang X., Mo H. J., van den Bosch F. C., 2008, *ApJ*, 676, 248
- Yang X., Mo H. J., van den Bosch F. C., Pasquali A., Li C., Barden M., 2007, *ApJ*, 671, 153
- Zwaan M. A., Meyer M. J., Staveley-Smith L., Webster R. L., 2005, *MNRAS*, 359, L30